

SOLIDserver Administrator Guide

Version 7.1

SOLIDserver Administrator Guide

SOLIDserver Administrator Guide

Revision: #87487

Publication date June 28, 2019

Copyright © 2000-2019 EfficientIP

All specifications and information regarding the products in this document are subject to change without notice and should not be construed as a commitment by EfficientIP. EfficientIP assumes no responsibility or liability for any mistakes or inaccuracies that may appear in this document. All statements and recommendations in this document are believed to be accurate but are presented without warranty. Users must take full responsibility for their application of any product.

Table of Contents

About This Guide	xviii
Documentation Organization	xviii
Documentation Convention	xix
I. Starting	1
1. Hardware Appliance Specificities	2
SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000	2
SDS-260	3
SDS-550, SDS-1100 and SDS-2200	4
SDS-3300 and SDS-Blast Series	6
2. First Installation and Basic Configuration	8
Prerequisites	8
SDS-50 First Installation	8
SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000 First Installation	10
SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast Series First Installation	11
SOLIDserver Software First Installation from a USB Flash Drive	17
Basic Network Configuration	18
3. Using SOLIDserver for the First Time	24
Connecting to SOLIDserver	24
Requesting and Adding a License	25
Defining the Internal Module Setup	26
Configuring SOLIDserver Network and Services	27
Creating SOLIDserver Users	27
4. Understanding the GUI	28
SOLIDserver Main Dashboard	29
Sidebar	29
Top Bar	35
Breadcrumb	38
Menu	40
Account Configuration	43
Listing Page	44
Properties Page	60
Charts	61
Wizards	65
Quick Wizards	71
Bookmarks	74
II. Configuring SOLIDserver	78
5. Configuring the Time and Date	79
Configuring NTP Servers	79
Forcing the NTP Update	80
Setting the Appliance Time and Date Manually	81
6. Configuring the Network	82
Configuring Basic IP Addressing on an Interface	83
Setting the Routing	84
Setting the Hostname	86
Setting the DNS Resolver	86
Setting the Firewall	87
Setting up a VLAN Interface	90
Setting up an Ethernet Port Failover	92
Configuring a VIP	93
Setting up a VIF	96

Configuring a Media Interface	97
7. Configuring the Services	99
Handling Services	100
Configuring the SSH Remote Account	101
Changing the SFTP/SCP/RSYNC User Account Password	102
Managing the TFTP Upload Authorizations	103
Configuring the SMTP Relay	103
Configuring the HTTPS Certificate	104
Configuring DNS Guardian	105
Configuring GSLB Server	105
Managing the SNMP Service	106
Downloading the DNS/DHCP/DHCPv6 Configuration File	108
8. Managing the Licenses	110
Renewing a License	110
Deleting a License	113
III. Imports and Exports	115
9. Importing Data from a CSV File	116
The Import Wizard	117
Importing Data to the IPAM	118
Importing Data to the DHCP	134
Importing Data to the DNS	143
Importing Data to NetChange	146
Importing Data to Device Manager	147
Importing Data to VLAN Manager	150
Importing Data to VRF	153
Importing Data to SPX	156
Importing Data to the Administration Module	159
Managing Import Templates	163
10. Importing IPAM Data	164
VitalQIP data	164
Nortel NetID data	165
11. Importing DHCP Data	167
ISC configuration files	167
Alcatel-Lucent VitalQIP configuration files	168
Microsoft configuration files	170
Infoblox configuration files	171
MetalP configuration files	172
Nortel NetID configuration files	173
12. Importing DNS Data	175
BIND archive files	175
VitalQIP archive files	176
13. Exporting Data	178
The Export Wizard	179
Browsing the Exports Database	180
Configuring Exports	180
Exporting Data To Reimport It Later	182
Managing Scheduled Exports	186
Managing Scheduled Exports Configuration Files	187
Managing Export Templates	187
IV. Dashboards	189
14. Building Dashboards	190
Browsing the Dashboards	190
Assigning a Gadget from the Dashboard	191
Organizing Gadgets on a Dashboard	191

Hiding Gadgets from a Dashboard	192
15. Managing Gadgets	193
Browsing Gadgets	193
Understanding the Gadget Types	195
Gadgets Displayed by Default	198
Creating Gadgets	207
Assigning Gadgets to a Dashboard	211
Displaying or Hiding a Gadget from the Page My Gadgets	213
Editing Gadgets	214
Granting Other Users Access to the Gadgets	216
Enabling or Disabling Gadgets	218
Deleting Gadgets	219
V. IPAM	220
16. Managing Spaces	222
Browsing Spaces	222
Adding a Space	223
Editing a Space	223
Deleting a Space	224
Defining a Space as a Group Resource	224
17. Managing Networks	225
Browsing Networks	225
Adding Networks	227
Editing Networks	234
Splitting Networks	235
Merging Networks	236
Moving Networks	237
Discovering the Assigned IP Addresses in a Network	237
Using Network Map to Display Assigned IP Addresses	238
Managing or Unmanaging Networks	239
Automating the DNS Records Update	240
Automating the DHCP Statics Reservation	240
Associating Networks With a VLAN	241
Deleting Networks	241
Defining a Network as a Group Resource	241
18. Managing Pools	243
Browsing Pools	243
Adding a Pool	244
Reserving a Pool	245
Resizing a Pool	245
Deleting a Pool	246
Defining a Pool as a Group Resource	246
19. Managing IP Addresses	247
Browsing IP Addresses	247
Adding an IP Address	249
Editing an IP Address	253
Configuring and Managing IP Address Aliases	254
Configuring Multiple A Records for an IP Address	257
Renaming IPv4 Addresses Massively	259
Moving IP Addresses	260
Pinging an IP Address	262
Deleting an IP Address	262
Restoring an IP Address	263
Automating the DNS Records Update From the Page All addresses	263
Automating the DHCP Statics Reservation From the Page All addresses	264

Updating Device Manager From the Page All addresses	264
20. Setting Up a Transition From IPv4 to IPv6	265
Transition Specificities	265
Limitations	265
Configuring the IPv4 to IPv6 Transition	266
Activating the IPv4 to IPv6 Transition	267
21. Managing IPAM Templates	271
Creating Template Classes in Class Studio	271
Creating IPAM Templates	272
Applying a Template	278
22. Using VLSM to Manage Your IPAM Network	280
Choosing the Method That Suits Your Needs	280
Managing a Space-Based VLSM Organization	284
Managing a Network-Based VLSM Organization	294
Using the VLSM Hierarchy to Delegate Management	300
VI. DHCP	303
23. Deploying DHCP Smart Architectures	305
Implementing DHCP Smart Architectures	306
DHCP Vendors Compatible with Smart Architectures	309
Building a Highly Available Service With Smart Architectures	310
24. Managing DHCP Smart Architectures	311
Browsing DHCP Smart Architectures	311
Adding a DHCPv4 Smart Architecture	312
Adding a DHCPv6 Smart Architecture	319
Editing a DHCP Smart Architecture	324
Handling the Status Locked Synchronization	329
Deleting a DHCP Smart Architecture	330
Defining a DHCP Smart Architecture as a Group Resource	331
25. Managing DHCP Servers	332
Browsing DHCP Servers	332
Managing EfficientIP DHCP Servers	334
Managing Agentless Microsoft DHCP Servers	337
Managing ISC DHCP Servers	342
Editing a DHCP Server	351
Configuring DHCP Options at Server Level	352
Deleting a DHCP Server	354
Defining a DHCP Server as a Group Resource	354
26. Managing DHCP Scopes	355
Browsing Scopes	355
Adding a DHCP Scope	356
Editing a DHCP Scope	357
Configuring DHCP Options at Scope Level	358
Defining a Specific IPAM Space for a Scope	359
Defining a Specific Failover Channel for a Scope	360
Replicating Scope Data in the IPAM	361
Configuring Multiple Scopes for a Network Segment	362
Copying or Moving DHCPv4 Scopes	362
Deleting a DHCP Scope	363
DHCP Relay Agents	364
Defining a DHCP Scope as a Group Resource	364
27. Managing Fixed Reservations	366
statics	366
groups	376
28. Managing Dynamic Addressing	379

ranges	379
leases	388
Restricting Access	396
Configuring the PXE	399
Preventing IP Address Duplication	401
29. Managing Failover Channels	403
DHCP Failover Principles and Operational States	403
Browsing the DHCP Failover Channels Database	406
Getting Familiar with the Failover Channels' Columns	406
Switching a DHCP server to Partner-down	408
30. Configuring DHCP Options	410
Setting DHCP Options	411
Customizing DHCP Options	412
DHCP Vendor Class Identifier	415
Option 82: Relay Agent Information	415
Option 43: Vendor Specific Information	417
31. Configuring DHCPv6 Prefix Delegation	420
Prerequisites	420
Specificities	420
Limitations	420
Browsing the DHCPv6 Prefix Delegations	420
Adding DHCPv6 Prefix Delegations	421
Deleting DHCPv6 Prefix Delegations	422
32. Monitoring and Reporting DHCP Data	423
Monitoring DHCP Servers From their Properties Page	423
Monitoring DHCP Servers From the Page Analytics	424
Monitoring DHCP Servers Using Rules	430
Generating DHCP Reports	432
VII. DNS	433
33. Deploying DNS Smart Architectures	435
Master/Slave Smart Architecture	435
Multi-Master Smart Architecture	436
Stealth Smart Architecture	437
Single-Server Smart Architecture	437
Farm Smart Architecture	438
34. Managing DNS Smart Architectures	439
Browsing DNS Smart Architectures	439
Adding a DNS Smart Architecture	440
Editing a DNS Smart Architecture	451
Handling the Status Locked Synchronization	456
Deleting a DNS Smart Architecture	457
Defining a DNS Smart Architecture as a Group Resource	458
35. Managing DNS Servers	459
Browsing DNS Servers	460
Managing EfficientIP DNS Servers	461
Managing Agentless Microsoft DNS Servers	465
Managing BIND DNS Servers	468
Managing Generic DNS Servers	478
Managing Nominum ANS Servers	480
Managing Amazon Route 53 Servers	482
Synchronizing DNS Servers	490
Editing DNS Servers	490
Securing the Management of DNS Servers Using a TSIG Key	491
Deleting DNS Servers	492

Defining a DNS Server as a Group Resource	493
36. Configuring DNS Servers	494
Configuring DNS Forwarding at Server Level	494
Configuring DNS Recursion at Server Level	496
Configuring DNS Notify Messages at Server Level	499
Restricting DNS Queries at Server Level	501
Limiting Zone Transfers at Server Level	503
Configuring a Blackhole	504
Configuring Client Resolver Cache Options at Server Level	505
Configuring EDNS Options at Server Level	506
Improving the Server Performance	507
Configuring a Sortlist at Server Level	507
Limiting the Number of Responses of a Server	508
Configuring DNS64	510
Configuring DNS Sources at Server Level	514
dynamic update on a server	517
Configuring Access Control Lists For a Server	519
Configuring DNS Keys	521
Including Non-Supported DNS Settings	523
Configuring Anycast DNS	523
Integrating Cisco Umbrella	535
37. Managing DNS Views	538
Browsing DNS Views	538
Adding DNS Views	539
Editing DNS Views	542
Editing the Order of the Views	543
Deleting DNS Views	544
Defining a DNS View as a Group Resource	544
Going Back to Managing Zones Without Views	545
38. Configuring DNS Views	546
Configuring DNS Forwarding at View Level	546
Configuring DNS Notify Messages at View Level	548
Configuring DNS Recursion at View Level	549
Restricting DNS Queries at View Level	551
Limiting Zone Transfer at View Level	554
Configuring Client Resolver Cache Options at View Level	555
Configuring EDNS Options at View Level	556
Configuring a Sortlist at View Level	556
Configuring DNS Sources at View Level	557
39. Managing DNS Zones	561
Browsing DNS Zones	561
Adding DNS Zones	563
Synchronizing DNS Zones	578
Editing DNS Zones	579
Converting DNS Zones	585
Adding or Removing an NS Record	586
Copying or Moving DNS Zones	586
Setting Properties on Multiple DNS Zones	587
Forcing the Update of DNS Zones' Content	593
Disabling and Enabling DNS Zones	594
Deleting DNS Zones	595
Defining a DNS Zone as a Group Resource	595
40. Configuring DNS Zones	596
Managing DNS Zone Delegation	596

Configuring DNS Forwarding at Zone Level	598
Configuring DNS Notify Messages at Zone Level	599
Managing DNS Security	602
Configuring DNS Sources at Zone Level	606
41. Managing DNS Resource Records	610
Browsing DNS Resource Records	612
Adding Resource Records	612
Editing Resource Records	628
Configuring the Delegation at Record Level	630
Copying or Moving Records	631
Changing the Hostname Convention	632
Deleting Resource Records	633
Load Balancing with Round-Robin	633
SPF Record	634
42. Implementing Dynamic Update	635
Configuring Dynamic Update	635
Configuring Secure Dynamic Update	636
Disabling Dynamic Update and Deleting Keys	643
43. DNS Firewall (RPZ)	646
Browsing RPZ Zones	647
Adding RPZ Zones	648
Editing RPZ Zones	649
Ordering RPZ Zones	651
Converting RPZ Zones	651
Deleting RPZ Zones	652
Managing RPZ Rules	652
Overriding RPZ Rules	664
RPZ Limitations	666
44. Hybrid DNS Service	667
Checking the Compatibility with Hybrid	667
Switching to Hybrid DNS	670
Forcing Compatibility with Hybrid	673
Switching Back to BIND	674
Administrating the Backup and Restoration of Hybrid Configurations	675
45. DNSSEC	676
Managing DNSSEC on Authoritative Servers	677
Managing DNSSEC on Recursive Servers	694
46. HSM	700
Prerequisites	700
Limitations	701
Browsing HSM Servers	701
Setting Up the HSM	703
Best Practices to Stop Using the HSM	709
47. Monitoring and Reporting DNS Data	712
Monitoring DNS Servers From their Properties Page	712
Monitoring DNS Servers From the Page Analytics	713
Monitoring DNS Queries and Answers	719
Generating DNS Reports	721
VIII. Global Policies	723
48. Inheritance and Propagation	724
Prerequisites	725
Limitations	726
Configuring the Inheritance of a Parameter Value	726
Configuring the Propagation of a Parameter Value	728

Setting Class Parameters	729
Reconciling Class Parameters	730
49. Managing Advanced Properties	732
Prerequisites	733
Browsing Advanced Properties	733
Selecting the Advanced Properties Displayed by Default	734
Configuring IPAM Advanced Properties	735
Configuring DHCP Advanced Properties	746
Configuring DNS Advanced Properties	748
Setting Advanced Properties	750
IX. Application	752
50. Configuring Application	753
Prerequisites	753
Limitations	753
Configuring and Enabling the Service GSLB Server	754
51. Managing Applications	756
Browsing Applications	756
Adding and Deploying Applications	757
Adding and Deploying Applications and Traffic Policies	759
Editing Applications	761
Deleting Applications	762
52. Managing Pools	763
Browsing Pools	763
Adding Pools	764
Editing Pools	766
Deleting Pools	766
53. Managing Nodes	768
Browsing Nodes	768
Adding Nodes	770
Editing Nodes	773
Managing or Unmanaging Nodes	774
Deleting Nodes	774
X. Guardian	775
54. Configuring Guardian	776
Prerequisites	776
Limitations	776
Enabling the Service DNS Guardian	777
55. Managing Guardian Configuration	779
Browsing Guardian Configuration	779
Editing Guardian Configuration	782
Configuring Guardian on Recursive and Authoritative Server	783
56. Managing Guardian Cache	784
Displaying Guardian Cache Content	784
Resetting Guardian Cache	786
Saving Guardian Cache	787
Restoring Guardian Cache	787
Forcing Cache Entries as Expired	788
Clearing Guardian Cache Manually	788
Clearing Guardian Cache Automatically	789
Sharing the Cache between Several Guardian servers	790
57. Managing Guardian Statistics	793
Managing Guardian server Statistics	793
Managing Guardian Client Statistics	802
Monitoring Guardian from the GUI	807

58. Managing Guardian Protection	820
Enabling Guardian Protection	820
Managing Guardian Rescue Mode	821
Managing Guardian Lists	824
Managing Guardian Views	829
Managing Guardian Policies	833
Managing Triggers	836
Disabling Guardian Protection	849
XI. NetChange	850
59. Managing Network Devices	852
Browsing Network Devices	852
Adding Network Devices	854
Importing Network Devices Using Discovery Protocols	855
Enabling or Disabling the 802.1X Authentication Protocol	857
Editing the SNMP Properties of a Network Device	858
Refreshing the Network Devices Database	859
Connecting to a Network Device Via a Console	861
Making a Network Device Snapshot	861
Creating Network Devices in Device Manager	862
Deleting Network Devices	862
Defining a Network Device as a Group Resource	862
60. Managing Routes	863
Prerequisites	863
Limitations	863
Browsing Routes	863
Enabling the Registry Key Required to Display the VRF Routes	865
61. Managing VLANs	867
Browsing VLANs	867
Adding a VLAN	868
Editing a VLAN	868
Deleting a VLAN	869
62. Managing Ports	870
Browsing Ports	871
Enabling or Disabling a Port	872
Editing a Port Interconnection	872
Editing a Port Speed and Duplex Mode	873
Updating a Port Description	874
Managing the 802.1X Authentication Protocol on a Port	874
Restricting Access to a Port with the Protocol Port-security	875
Limiting Port Edition Rights to Specific Groups of Users	878
Configuring VLAN Tagging on a Port	880
Refreshing the Ports Database	883
63. Managing Configuration Versioning	884
Prerequisites	884
Limitations	884
Browsing Configuration Files	884
Managing Connection Profiles	886
Configuring the Versioning of a Network Device	888
Refreshing a Configuration File	890
Comparing Configuration Files	891
Monitoring the Configuration Versioning in the Logs	892
Downloading Versioning Information	893
Configuring Advanced Versioning Options	894
Disabling Versioning on a Network Device	895

64. Managing Addresses	897
Browsing Addresses	897
65. Managing Discovered Items	899
Browsing Discovered Items	899
Refreshing the Discovered Items Database	901
Populating Device Manager	901
Creating the IP Address of a Discovered Item in the IPAM	902
Purging the Discovered Items Database	902
Tracking the Discovered Items History of a Specific Device	903
66. Managing Statistics	904
Displaying NetChange Statistics	904
Displaying Network Device Statistics	904
Displaying Port Statistics	904
67. Monitoring and Tuning	906
Generating NetChange Reports	906
Keeping NetChange Data Up-to-date	906
Synchronizing the Network Devices with a CSV File	907
Customizing the Type of Network Devices	908
XII. Workflow	910
68. Granting Access to Workflow Classes	911
69. Managing Outgoing Requests	912
Browsing Outgoing Requests	912
Adding Requests for Creation	913
Adding Requests for Edition	914
Adding Requests for Deletion	915
Editing Requests	917
Canceling Requests	919
70. Managing Incoming Requests	920
Browsing Incoming Requests	920
Managing the Requests Content	920
Administrating Requests Using the Default Statuses and Options	921
Administrating Requests Using Your Own Settings	924
71. Executing Requests	925
Executing Requests Using the Execute Option	925
Executing Requests Using Classes	926
72. Customizing the Requests Administration	930
Editing the Workflow Statuses	931
Editing the Email Notification Details	933
Adding a Workflow Status	934
Customized Statuses Best Practices	935
XIII. Device Manager	937
73. Managing Devices	938
Browsing Devices	938
Adding Devices	939
Duplicating Devices	946
Merging Devices	946
Deleting Devices	947
74. Managing Ports and Interfaces	948
Browsing Ports and Interfaces	948
Adding Ports and Interfaces	949
Editing Ports and Interfaces Properties	956
Tracking Changes on the Page All ports & interfaces	960
Deleting Ports and Interfaces	961
75. Managing the Interaction with the IPAM	963

Assigning IP Addresses to an Interface Using their MAC Address	963
Managing the IP Addresses / Interfaces Link from the IPAM	966
Editing the Devices Topology from the IPAM	969
76. Rules Impacting Device Manager	971
DHCP Rules Impacting Device Manager	971
Adding Device Manager Rules	971
Enabling or Disabling Device Manager Rules	972
XIV. VLAN Manager	973
77. Managing VLAN Domains	974
Browsing VLAN Domains	974
Adding VLAN Domains	974
Editing VLAN Domains	975
Deleting VLAN Domains	976
Defining a VLAN Domain as a Group Resource	976
78. Managing VLAN Ranges	977
Browsing VLAN Ranges	977
Adding VLAN Ranges	978
Editing VLAN Ranges	978
Deleting VLAN Ranges	980
Defining a VLAN Range as a Group Resource	980
79. Managing VLANs	981
Browsing VLANs	981
Adding VLANs	982
Editing VLANs	983
Deleting VLANs	983
80. Managing the IPAM/VLAN Interaction	984
Displaying the IPAM/VLAN Interaction Advanced Properties	984
Configuring the IPAM/VLAN Interaction	985
Removing the IPAM/VLAN Interaction	986
XV. VRF	988
81. Managing Virtual Routing and Forwarding	989
Browsing VRFs	989
Adding VRFs	989
Editing VRFs	990
Deleting VRFs	990
82. Managing VRF Route Targets	992
Browsing VRF Route Targets	992
Adding VRF Route Targets	992
Deleting VRF Route Targets	993
XVI. SPX	995
83. Configuring SPX	996
Enabling the SPX Classes	996
Enabling the SPX Rules	996
Configuring the SPX Connection	996
Editing the Connection to the RIPE or APNIC	999
84. Managing SPX Persons	1002
Browsing SPX Persons	1002
Adding SPX Persons	1002
Editing SPX Persons	1003
Registering SPX Person Changes	1004
Deleting SPX Persons	1004
85. Managing SPX Networks	1006
Browsing SPX Networks	1006
Adding SPX Networks	1007

Editing SPX Networks	1010
Registering SPX Network Changes	1011
Validating a New Assignment Window	1012
Deleting SPX Networks	1012
86. Managing SPX AS Numbers	1014
Browsing SPX AS Numbers	1014
Adding SPX AS Numbers	1015
Editing SPX AS Numbers	1016
Deleting SPX AS Numbers	1016
XVII. Rights Management	1018
87. Managing Groups	1019
Browsing Groups of Users	1019
Adding Groups of Users	1020
Managing the Resources of a Group of Users	1021
Managing the Rights of a Group of Users	1025
Managing the Users of a Group of Users	1027
Editing Groups of Users	1028
Enabling or Disabling Groups of Users	1028
Deleting Groups of Users	1029
88. Managing Users	1030
Browsing Users	1030
Adding Users	1031
Editing Users	1032
Changing the User Password	1033
Configuring the User Password Complexity	1034
Configuring Users Connection Parameters	1034
Configuring User Sessions	1036
Enabling or Disabling Users	1037
Generating User Reports	1037
Deleting Users	1037
89. Managing Authentication Rules	1038
Browsing Authentication Rules	1039
Adding Authentication Rules	1039
Editing Authentication Rules	1045
Enabling or Disabling Authentication Rules	1045
Deleting Authentication Rules	1045
XVIII. Administration	1047
90. Centralized Management	1048
Browsing the Centralized Management Database	1048
Configuring SOLIDserver to Remotely Manage Other Appliances	1051
Adding Remote Appliances	1052
Managing the Services and Network Configuration of Another Appliance	1053
Monitoring the Appliances Managed from the Centralized Management	1054
Configuring Two Appliances in High Availability	1058
Editing Remote Appliances	1061
Managing a High Availability Configuration	1061
Replacing Appliances Managed Remotely	1074
Deleting Remote Appliances	1075
91. Monitoring	1078
Managing Reports	1078
Managing Alerts	1088
Managing the Logs	1095
Monitoring the Appliance Statistics	1097
Tracking Sessions	1099

Tracking Users	1099
Managing SNMP Profiles	1101
Monitoring Using SNMP	1103
Displaying Netstat Data	1104
Sizing the Database Tables	1104
92. Maintenance	1106
Managing the HTTPS Certificate	1106
Managing Files from the Local Files Listing	1112
Using the Maintenance mode	1119
Updating the Macros and Rules	1120
Clearing the Appliance Cache	1120
Troubleshooting	1120
Managing Backups and Restoring Configurations	1123
Shutting Down and Rebooting	1129
93. Upgrading	1133
Prerequisites	1133
Upgrading an Appliance	1133
Upgrading Appliances Managed Remotely	1135
Upgrading Appliances in High Availability	1136
Troubleshooting the Upgrade	1142
XIX. Customization	1150
94. Customizing the GUI	1151
Customizing SOLIDserver Login Page With an Image	1151
Customizing SOLIDserver Home Page Welcome Banner	1153
Customizing the Interface Names and Fields	1155
95. Managing Smart Folders	1157
Browsing Smart Folders	1157
Adding Smart Folders	1158
Editing Smart Folders	1159
Sharing Smart Folders	1159
Deleting Smart Folders	1160
96. Managing IPv6 Labels	1161
Limitations	1161
Adding Labels	1161
Displaying or Hiding Labels	1162
Editing Labels	1162
Deleting Labels	1163
97. Configuring Classes	1164
Browsing Class Studio Database	1165
Managing Classes	1167
Configuring the Classes Content	1174
Class Studio Syntax	1221
Defining a Class as a Group Resource	1222
98. Configuring Custom Databases	1223
Browsing Custom DB	1223
Adding a Custom DB	1224
Editing a Custom DB	1224
Deleting a Custom DB	1225
Configuring a Custom DB with Custom Data	1225
99. Managing Customization Packages	1227
Browsing the Packages Database	1227
Uploading Packages	1228
Creating Packages	1228
Editing Packages	1230

Installing Packages	1230
Uninstalling Packages	1231
Downloading Packages	1231
Deleting Packages	1231
A. Matrices of Network Flows	1233
SOLIDserver	1234
IPAM	1235
DHCP	1236
DNS	1237
NetChange	1239
Remote Management	1240
B. Multi-Status Messages	1241
DHCP Multi-Status Messages	1241
DNS Multi-Status Messages	1241
C. Default Gadgets	1243
D. DHCP Options	1245
Basic Options	1245
Server Parameters	1246
Lease Information Options	1247
WINS/NetBIOS Options	1247
Host IP Options	1247
Interface Options	1248
Servers Options	1249
BOOTP Compatibility Options	1250
DHCP Packet Fields Options	1251
Microsoft DHCP Client Options	1252
NetWare Client Options	1252
NIS/NISplus Options	1253
Miscellaneous	1254
Vendor MSFT Options	1254
Vendor Nwip Options	1254
E. MAC Address Types References	1255
F. DNS Resource Records Related Fields	1257
G. Advanced Properties	1261
H. User Tracking Services Filter	1262
I. SNMP Metrics	1266
Prerequisites	1266
Understanding the SNMP Metrics Presentation	1267
Retrieving SNMP Metrics via CLI	1267
Monitoring the Hardware	1268
Monitoring the System	1270
Monitoring the DHCP Service	1274
Monitoring the DNS Service	1275
Monitoring DNS Guardian	1278
J. Class Studio Pre-defined Variables	1286
K. Configuring RADIUS	1289
Configuring FreeRADIUS	1289
Configuring RADIUS with Cisco ACS	1290
Configuring OneTime Password with Token Authentication	1291
L. Using Remote Authentication for SSH Connections to SOLIDserver	1294
Configuring LDAP Authentication for SSH Connections	1294
Configuring RADIUS Authentication for SSH Connections	1299
M. Configuring Non-Supported Options	1303
Prerequisites	1304

Limitations	1304
Configuring Non-Supported Firewall Rules	1305
Configuring Non-Supported Apache Settings	1307
Configuring Non-Supported Unbound Settings	1309
Configuring Non-Supported NSD Settings	1311
Configuring Non-Supported BIND Settings	1313
Configuring Non-Supported SNMP Settings	1317
Configuring Non-Supported DHCP Configurations	1318
Configuring Non-Supported NTP Settings	1320
Configuring Non-Supported syslog-ng Settings	1321
Configuring Non-Supported PostgreSQL Settings	1322
Index	1323

About This Guide

SOLIDserver is a hardware or software appliance suite that allows to manage from one device a network at all levels, from the IP addresses to the network devices, through key services, systems and protocols.

The Administrator Guide describes and details the modules you might have purchased with your license. This guide does not detail the existing license types and what they may contain or lack. Note that some configurations described in this document should not be handled by end users if they do not have previous knowledge of the basic principles of certain protocols and what the operations imply on the network configuration.

Documentation Organization

The guide is divided into the following **parts**:

- **Starting**: the hardware appliances description, the installation of SOLIDserver on hardware and software, the basic network configuration, the appliance first use procedures, a comprehensive presentation of the GUI and how to use gadgets to build your modules' dashboard.
- **Configuring SOLIDserver**: the system configuration possibilities of an appliance such as the network configuration or the services.
- **Imports and Exports**: all the data import and export options available.
- **Dashboards**: all the options available in the Dashboards module.
- **IPAM**: all the options available to assist you in IP addresses management.
- **DHCP**: all the options available in the DHCP protocol dedicated module.
- **DNS**: all the options available in the DNS protocol dedicated module.
- **Global Policies**: the options available in all modules: parameters inheritance and propagation, advanced properties, imports, exports, reports, alerts and Smart Folders.
- **Application**: all the options available in the module Application.
- **Guardian**: all the options available in the module Guardian.
- **NetChange**: all the options available in the module NetChange, dedicated to devices connected to your network.
- **Workflow**: all the options available in the Workflow dedicated module.
- **Device Manager**: all the options available in the module Device Manager dedicated to devices, ports and interfaces personalized management.
- **VLAN Manager**: all the options available in the module VLAN Manager dedicated to Virtual Local Area Network personalized management.
- **VRF**: all the options in the module VRF dedicated to Virtual Routing and Forwarding.
- **SPX**: all the options available in the module SPX dedicated to the Service Provider eXtension module that allows you to manage your RIPE or APNIC database from SOLIDserver.
- **Rights Management**: all the options available to manage users, groups of users, their authentication and the operations they can perform.
- **Administration**: all the options and operations available in the module Administration module from configuring appliances in high availability or remotely managing them, to monitoring, maintaining and upgrading the appliance.

- **Customization**: a description of all the options available in the Administration module to customize your appliance: from images, IPv6 labels and Smart Folders to custom databases, classes and customization packages.

At the end of the guide, you can also find **appendices** containing further details regarding:

- **Matrices of Network Flows** details the network flows of the DNS, the DHCP, the IPAM as well as the High Availability Management or NetChange.
- **Multi-Status Messages** contains the existing messages returned by the column Multi-status.
- **Default Gadgets** describes the gadgets available by default: their type, purpose and the dashboard they are displayed when relevant.
- **DHCP Options** includes options and parameters from basic options to lease information, host IP, interfaces, servers, BOOTP, Microsoft, NetWare NIS/NISplus or even vendors options.
- **MAC Address Types References** displays the reference number, in the GUI, of DHCP statics supported MAC types.
- **DNS Resource Records Related Fields** displays the fields that need to be configured when adding resource record to a zone.
- **Advanced Properties** provides a graphic showing all the advanced properties.
- **User Tracking Services Filter** details the use of the filters available on the page User tracking.
- **SNMP Metrics** provides a list of the most relevant SOLIDserver indicators you can monitor through an external solution.
- **Class Studio Pre-defined Variables**: provides a table detailing the values available when configuring a pre-defined variable class object.
- **Configuring RADIUS**: provides procedures to configure FreeRadius and RADIUS with Cisco ACS and make them compatible with SOLIDserver.
- **Using Remote Authentication for SSH Connections to SOLIDserver**: provides the configuration details to use LDAP or RADIUS authentication and grant access to existing LDAP/RADIUS users to SOLIDserver via SSH.
- **Configuring Non-Supported Options**: provides advanced configuration details to incorporate non-supported firewall rules and options for the services Apache, Unbound, NSD, BIND, SNMP, DHCP, NTP, syslog-ng and PostgreSQL.

Documentation Convention

Each part of this guide is divided into **chapters** where the available operations are detailed in **procedures**. Throughout the guide, you will find the following elements.

Element	Description
Procedure	All the procedures detail step by step the available configurations and possible operations. They contain highlighted words and icons matching the Graphical User Interface (GUI).
Name	All key words to browse the GUI: page name, wizard title, column name...
BUTTON	The buttons in the GUI: OK, EDIT, DELETE, CANCEL, UPDATE, etc.
Menu	The menus and their entries. The navigation within menu entries is symbolized by arrows as such: menu > option > sub-option .

Part I. Starting

This part details the very first operations to connect to SOLIDserver, for hardware and software appliances as well as a presentation of the GUI and the modules' dashboard. It contains the following chapters:

- [Hardware Appliance Specificities](#): describes our hardware appliance suite. Depending on the model you chose, the front panel and available buttons and possible configurations from the hardware appliance differ.
 - [First Installation and Basic Configuration](#): describes the hardware appliance installation procedures for all hardware models, the installation via USB and the basic network configuration to complete the appliance configuration or to reset the default configuration details.
 - [Using SOLIDserver for the First Time](#): describes the first steps and best practices to follow after setting SOLIDserver with an IP address: from logging in and configuring the appliance service to creating users.
 - [Understanding the GUI](#): describes the default features of SOLIDserver Graphical User Interface. It includes a presentation of the navigation philosophy, some tips and all the extra functionalities we provide: bookmarks, quick wizards, global search, etc.
-

Chapter 1. Hardware Appliance Specificities

There are different hardware appliances models available:

1. The hardware appliance of SOLIDserver SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000.
2. The hardware appliance of SOLIDserver SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast series (SDS-4000, SDS-5000 and SDS-5500).

SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000

SOLIDserver SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000 front panel allows you to plug in one or more physical interfaces. It is also possible to plug in a console cable to visualize the display output on another computer/screen. Using the server panel, you can easily set up basic network configuration.

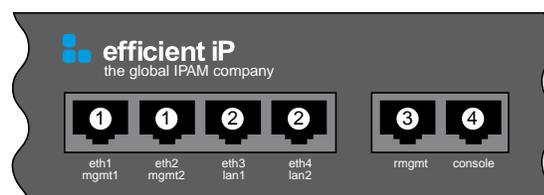


Figure 1.1. Front panel ports

- ❶ **eth1/eth2:** These ports allow you to plug in one or more cables for the management. To understand how SOLIDserver manages all physical interfaces, refer to the section [Configuring SOLIDserver](#).
- ❷ **eth3/eth4:** These ports allow you to plug in one or more cables for the network services. To understand how SOLIDserver manages all physical interfaces, refer to the section [Configuring SOLIDserver](#).
- ❸ **rmgmt:** This port is IPMI and LAN dedicated. The IPMI protocol leverages an out-of-band network (typically dedicated for server monitoring and management), that provides a flawless and secure path for mission-critical applications when regular in-band connectivity is lost or is unresponsive.
- ❹ **console:** This port allows you to plug in a console cable to visualize the output on a terminal device. The console port on SOLIDserver is an asynchronous serial port. The console port is configured as data terminal equipment (DTE). The console port uses RJ-45 connectors (Cisco). Adapters are available for connections to PC terminals, modems, or other external communication equipment. To connect a PC terminal to the console port, use either a RJ-45-to-RJ-45 rollover cable, a RJ-45-to-DB-25 female DTE adapter or the RJ-45-to-DB-9 female DTE adapter (labeled "TERMINAL"). The default parameters for the console port are:
 - 9600 baud.
 - 8 data bits.
 - No parity generated or checked.
 - 1 stop bit.
 - No Flow Control.

You may use the following software:

- Hyper terminal or Putty on Windows.
- minicom on Linux, MacOS, Unix.
- cu on Linux, MacOS, Unix.



Figure 1.2. Front panel LED, LCD screen and buttons

- 1 This LED glows green when the appliance is on.
- 2 This LED flashes red when the hard drive processes data.
- 3 This LED glows yellow in case of a new information event.
- 4 Button to start/stop the server. If you keep the button pushed, the server shuts down directly - this action is not recommended. If you push the button once, the server shuts down by itself.
- 5 Buttons to navigate through the LCD screen menus.

SDS-260

The sections below describe the hardware specificities of the SOLIDserver SDS-260 appliance.

SDS-260 Front Panel

SOLIDserver SDS-260 front panel allows you to set up the access IP address of the iDRAC via a keyboard/monitor. From the iDRAC interface, you can configure SOLIDserver IP address.

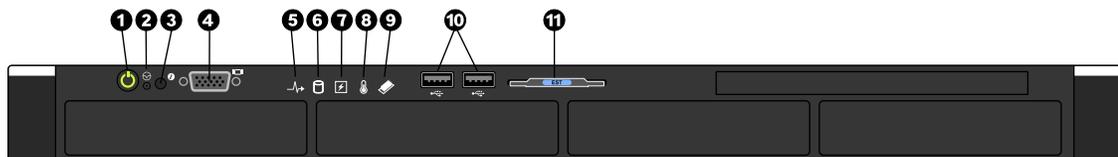


Figure 1.3. Front panel

- 1 The On/Off button. It is lit when the appliance is on.
- 2 The NMI button. It is inactive on SOLIDserver hardware appliances.
- 3 The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- 4 The VGA video connector allows you to connect a monitor to the system.
- 5 The health indicator. If the system is on and in good health, the indicator turns solid blue. The indicator flashes amber if the system is on or in standby, and if any error exists (for example, a failed fan or hard drive).
- 6 The hard drive indicator. The indicator flashes amber if there is a hard drive error.

- 7 The electrical indicator. The indicator flashes amber if the system experiences an electrical error (for example, voltage out of range, or a failed power supply unit or voltage regulator).
- 8 The temperature indicator. The indicator flashes amber if the system experiences a thermal error (for example, a temperature out of range or fan failure).
- 9 The memory indicator. The indicator flashes amber if a memory error occurs.
- 10 The USB ports 1 and 2. They are USB 2.0 compliant.
- 11 The information tag. Pull it out to see the hardware service tag and express service tag.

SDS-260 Back Panel

SOLIDserver SDS-260 back panel allows you to connect a power supply, your network devices and other types of equipment to your appliance.

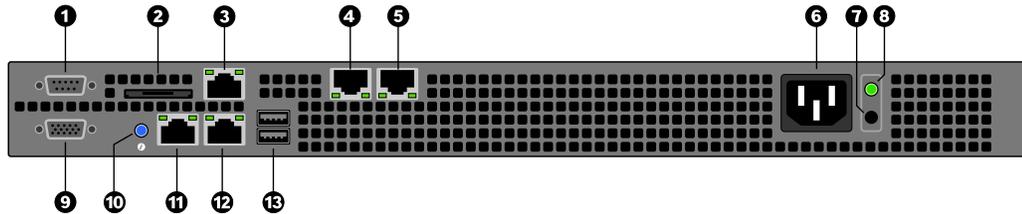


Figure 1.4. Back panel

- 1 The iDRAC serial port allows you to connect a serial device to the system.
- 2 The vFlash card slot.
- 3 The iDRAC port is dedicated to the remote management of the appliance.
- 4 The ethernet connector *bge1* allows you to connect network devices to your appliance.
- 5 The ethernet connector *bge0* serves the same purpose as *bge1*.
- 6 The power supply unit (PSU) allows you to plug the appliance to a power source.
- 7 The self-diagnostic button allows you to perform a quick health check on the cabled PSU of the system.
- 8 The AC power supply unit status indicator. When pressing the self-diagnostic button, it turns green if a valid power source is connected. Otherwise, the PSU is not connected or faulty.
- 9 The VGA video connector allows you to connect a monitor to the system.
- 10 The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- 11 The ethernet connector *bge2* allows you to connect network devices to your appliance.
- 12 The ethernet connector *bge3* serves the same purpose as *bge2*.
- 13 The USB ports 3 and 4. They are USB 3.0 compliant.

SDS-550, SDS-1100 and SDS-2200

The sections below describe the hardware specificities of the SOLIDserver SDS-550, SDS-1100 and SDS-2200 appliances.

SDS-550, SDS-1100, SDS-2200 Front Panel

SOLIDserver SDS-550, SDS-1100, SDS-2200 front panel allows you to set up the IP address of the iDRAC used to configure SOLIDserver, via an LCD screen or a keyboard/monitor. It also allows you to display the hardware information.

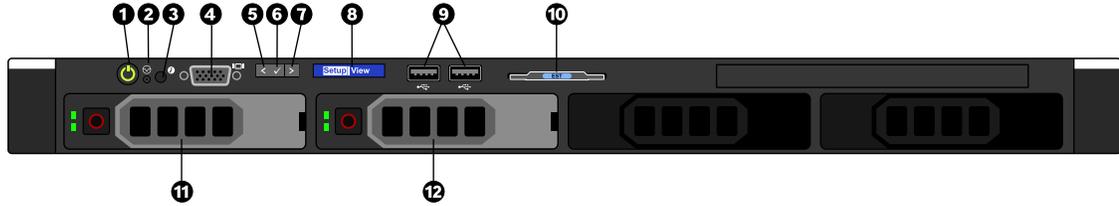


Figure 1.5. Front panel

- ❶ The On/Off button. It is lit when the appliance is on.
- ❷ The NMI button. It is inactive on SOLIDserver hardware appliances.
- ❸ The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- ❹ The VGA video connector allows you to connect a monitor to the system.
- ❺ The LCD menu left button. Moves the cursor back in one-step increments.
- ❻ The LCD menu selection button. Selects the menu item highlighted by the cursor.
- ❼ The LCD menu right button. Moves the cursor forward in one-step increments.
- ❽ The LCD screen used to set up the appliance iDRAC IP address or display hardware information
- ❾ The USB ports 1 and 2. They are USB 2.0 compliant.
- ❿ The information tag. Pull it out to see the hardware service tag and express service tag.
- ⓫ The hot swappable hard drive 1.
- ⓬ The hot swappable hard drive 2.

SDS-550, SDS-1100, SDS-2200 Back Panel

SOLIDserver SDS-550, SDS-1100 and SDS-2200 back panel allows you to connect a power supply, your network devices and other types of equipment to your appliance.

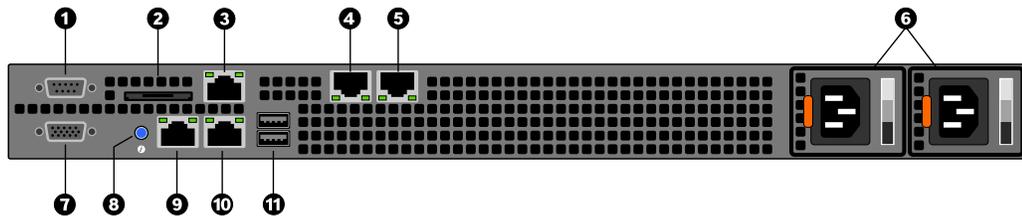


Figure 1.6. Back panel

- ❶ The iDRAC serial port allows you to connect a serial device to the system.
- ❷ The vFlash card slot.
- ❸ The iDRAC port. It is dedicated to the remote management of the appliance.

- ④ The ethernet connector *igb0* allows you to connect network devices to your appliance.
- ⑤ The ethernet connector *igb1* serves the same purpose as *igb0*.
- ⑥ The power supply unit (respectively PSU1 and PSU2) allow to plug the appliance to a power source. It is recommended to use both PSUs to prevent connection loss if one fails.
- ⑦ The VGA video connector allows you to connect a monitor to the system.
- ⑧ The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- ⑨ The ethernet connector *bge0* allows you to connect network devices to your appliance.
- ⑩ The ethernet connector *bge1* serves the same purpose as *bge0*.
- ⑪ The USB ports 3 and 4. They are USB 3.0 compliant.

SDS-3300 and SDS-Blast Series

The sections below describe the hardware specificities of the SDS-3300 and SDS-Blast Series (SDS-4000, SDS-5000 and SDS-5500) appliances.

SDS-3300 and SDS-Blast Series Front Panel

SOLIDserver SDS-3300 and SDS-Blast series (SDS-4000, SDS-5000 and SDS-5500) front panel allows you to set up the IP address of the iDRAC used to configure SOLIDserver, via an LCD screen or a keyboard/monitor. It also allows you to display the hardware information.

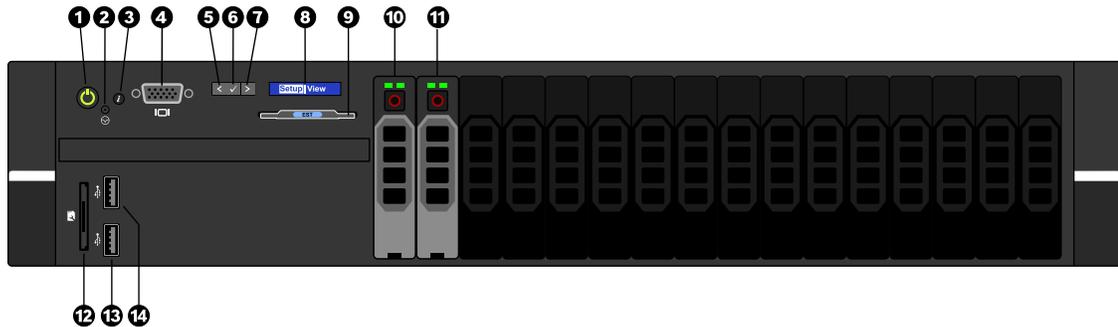


Figure 1.7. Front panel

- ① The On/Off button. It is lit when the appliance is on.
- ② The NMI button. It is inactive on SOLIDserver hardware appliances.
- ③ The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- ④ The VGA video connector allows you to connect a monitor to the system.
- ⑤ The LCD menu left button. Moves the cursor back in one-step increments.
- ⑥ The LCD menu selection button. Selects the menu item highlighted by the cursor.
- ⑦ The LCD menu right button. Moves the cursor forward in one-step increments.
- ⑧ The LCD screen used to set up the appliance iDRAC IP address or display hardware information.
- ⑨ The information tag. Pull it out to see the hardware service tag and express service tag.
- ⑩ The hot swappable hard drive 1.

- 11 The hot swappable hard drive 2.
- 12 The SD card plug. The card contains SOLIDserver installation image.
- 13 The USB port 1. It is USB 2.0 compliant.
- 14 The USB port 2. It is USB 2.0 compliant.

SDS-3300 and SDS-Blast Series Back Panel

SOLIDserver SDS-3300 back panel allows you to connect a power supply, your network devices and other types of equipment to your appliance. For SOLIDserver SDS-Blast Series, it also provides you with two 10 Gbps optical fiber slots, for a full DNS Guardian protection.

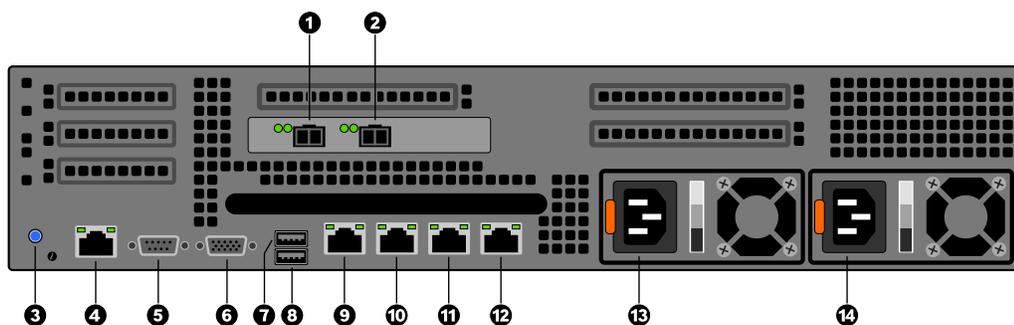


Figure 1.8. Back panel

- 1 The optical slot *ix1* allows you to connect a 10 Gbps LC optical connector¹. It is only available on SOLIDserver SDS-Blast Series appliances.
- 2 The optical slot *ix0* serves the same purpose as *ix1*.
- 3 The system identification button locates a particular appliance in a rack: push it to light it up both at the back and the front of the appliance. Push it again to turn off the lights.
- 4 The iDRAC enterprise port. It is dedicated to the remote management of the appliance.
- 5 The iDRAC serial port allows you to connect a serial device to the system.
- 6 The VGA video connector 1. It allows you to connect a monitor to the system.
- 7 The USB port 3. It is USB 3.0 compliant.
- 8 The USB port 4. It is USB 3.0 compliant.
- 9 The ethernet connector *igb0* allows you to connect network devices to your appliance.
- 10 The ethernet connector *igb1* serves the same purpose as *igb0*.
- 11 The ethernet connector *igb2* serves the same purpose as *igb0*.
- 12 The ethernet connector *igb3* serves the same purpose as *igb0*.
- 13 The power supply unit 1 (PSU1) allows you to plug the appliance to a power source. It is recommended to use both PSUs to prevent connection loss if one fails.
- 14 The power supply unit 2 (PSU2) serves the same purpose as PSU1.

¹We recommend using SFP+ Intel X520 adapters. For more details on the supported products, refer the Intel website at <https://www.intel.com/content/www/us/en/support/articles/000005528/network-and-i-o/ethernet-products.html>

Chapter 2. First Installation and Basic Configuration

This chapter is a recollection of the proper first installation and basic network configuration of SOLIDserver hardware appliances.

Following the installation procedures is required for the **first installation of hardware appliances** or if you reset it to its factory default settings. In either case, you must:

- Meet the [prerequisites](#).
- Follow the procedure of your hardware appliance model:
 - [SDS-50 First Installation](#)
 - [SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000 First Installation](#)
 - [SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast Series First Installation](#)
 - [SOLIDserver Software First Installation from a USB Flash Drive](#)

When the installation is complete, you can use a terminal or LCD screen for the appliance :

- Further the **basic network configuration** of the appliance hostname, interfaces, etc. via a terminal.
- Reset the default network settings from the LCD screen.

Fore more details, refer to the section [Basic Network Configuration](#).

Prerequisites

Before using SOLIDserver Graphical User Interface (GUI), you need to connect your device to the network and configure an IP address for it (i.e. the default gateway configured during the installation).

The management client is the computer from which you configure and manage SOLIDserver, make sure you meet the following requirements:

- You must have at least one of the following supported web browsers: *Google Chrome*, *Mozilla Firefox*, *Microsoft Internet Explorer*, *Microsoft Edge* or *Safari* (Mac only). We recommend using its latest and most stable version.
- You must disable any web browser pop-up blocker for the IP address and domain name configured for your appliance. Otherwise, you cannot manage SOLIDserver properly.

SDS-50 First Installation

SOLIDserver software is already installed on hardware appliances SDS-50, you need to configure the access IP address of your choice to manage it.

Prerequisites

- A SOLIDserver hardware appliances SDS-50.
- A computer with a serial port.

- A terminal emulator able to connect to a serial port installed on your computer.

Configuring SOLIDserver IP Address

To configure SOLIDserver and the IP address used to access it, you must use a terminal emulator.

To configure SOLIDserver on SDS-50 appliances

1. Plug in the appliance power cable.
2. Plug in the appliance to your computer serial port.
3. Plug in at least one physical interface Ethernet cable to the appliance.
4. Open a terminal emulator and start a session with a **Serial** connection. When the appliance has started, the page **WELCOME TO SOLIDSERVER** opens.
5. Log in using the login *root* and no password. The configuration window opens, it indicates the version: **SOLIDserver <version>**.
6. In the **Main Menu**, **S Quick Start** is highlighted.
7. Hit **Enter**. The page **Quick Start - Configure basic network settings** opens.
8. Hit **Enter** to confirm the quick start configuration. The menu **Physical interfaces** opens.
9. In the menu **Physical Interfaces**, select one or several interfaces. [*] indicates which interface(s) is selected. Use the keyboard arrows or the digits to highlight the interface of your choice and hit the **space** key to select or deselect an interface.

If you plugged in several physical interfaces Ethernet cables, you are configuring an Ethernet port failover.

10. Hit **Enter** to confirm the interface(s) selection. The menu **Network configuration** opens.
11. Configure the access to SOLIDserver:
 - Configure the interface IP address on the line **IP:** . Hit the down arrow to get to the next line.
 - Configure the interface netmask on the line **Netmask:** . Hit the down arrow to get to the next line.
 - Configure the interface default gateway on the line **Gateway:** . Hit the tabulation key to highlight OK.
12. Hit **Enter** to confirm your configuration and go back to the **Main Menu**.
13. The line **C Commit modifications to system** is highlighted.
14. Hit **Enter** to save the whole configuration. The window **Do you really want to commit modifications to system?** opens.
15. The button **Yes** is highlighted.
16. Hit **Enter** to commit your choice. When the configuration is saved, the window **Configuration applied** opens.
17. Hit **Enter** to close the window. The **Main Menu** is visible again.
18. Hit **E** to select **EXIT** and hit **Enter**. The configuration window closes.

Now your configuration is complete and you can connect to SOLIDserver software using the IP address you configured. For more details, refer to the chapter [Using SOLIDserver for the First Time](#).

SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000 First Installation

These types of SOLIDserver belong to the third generation hardware appliances, their front panel LCD screen can be used in two ways:

1. To set up the first network configuration.

Once the appliance has booted and is up and running, you have 30 seconds to press any arrow button to get into the edition mode. Through this mode you can edit the IP address as well as the netmask and the gateway. By default, the first physical interface is configured (*eth1* on the server panel, called *em0* in the system) with an IP address *192.168.1.1/255.255.255.0* and a gateway set to *0.0.0.0*. The configuration of the Basic Network using LCD display can be applied only on *eth1* with one IP address.

2. To visualize the network configuration.

The LCD screen displays at all times the following information when SOLIDserver is running: the hostname, serial number, IP address, prefix/netmask and gateway. During the very first configuration, if you let the 30 seconds timer run out, the default configuration is implemented and displayed on the screen after 90 seconds. No matter how you configured the appliance, the LCD screen provides a summary of the key network configuration data.

Prerequisites

- A SOLIDserver hardware appliance SDS-250, SDS-500, SDS-1000, SDS-2000 or SDS-3000.

Configuring SOLIDserver IP Address

From the front panel LCD screen, you can SOLIDserver and the IP address used to access it.

To configure SOLIDserver on SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000 appliances

1. Plug in an Ethernet cable on Eth1, the default interface configured on SOLIDserver.
2. Turn on SOLIDserver by pushing **⏻**. During the boot sequence, the LCD screen displays *Device booting*.
3. Once SOLIDserver is booted, a 30 seconds countdown is displayed. Press **⏻** to enter the setup.
4. Modify the IP address: to move from one octet to the other, press **⬅/➡**. To decrease/increase its value, press **⬇/⬆**. By default, the IP address is *192.168.1.1*.
5. To commit the new IP address, go to the last position of the menu (on the far right). A new Menu appears:
 - **⏻** Esc: Cancel the current modification and go back to the menu.
 - **⏻** No: Cancel the current modification and allows you to set a new value.
 - **⏻** Yes: Commit the value. The message *Performing* appears.

6. Modify the Netmask: to move from one octet to the other, press **◀/▶**. To decrease/increase its value, press **▼/▲**. By default, the netmask is `255.255.255.0`.
7. To commit the new Netmask, go on the last position of the menu (on the far right). A new Menu appears:
 - **◀** Esc: Cancel the current modification and go back to the menu.
 - **▼** No: Cancel the current modification and allows you to set a new value.
 - **▲** Yes: Commit the value. The message *Performing* appears.
8. Modify the Gateway: to move from one octet to the other, press **◀/▶**. To decrease/increase its value, press **▼/▲**. By default, the gateway is `0.0.0.0`.
9. To commit the new Gateway, go on the last position of the menu (on the far right). A new Menu appears:
 - **◀** Esc: Cancel the current modification and go back to the menu.
 - **▼** No: Cancel the current modification and allows you to set a new value.
 - **▲** Yes: Commit the value. The message *Performing* appears.
10. The basic network configuration is applied.

Now your configuration is complete and you can connect to SOLIDserver software using the IP address you configured. For more details, refer to the chapter [Using SOLIDserver for the First Time](#).

Keep in mind that the LCD screen can also be used to reset the hardware default configuration. For more details, refer to the section [Resetting the Default Network Settings From the LCD Screen](#).

SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast Series First Installation

These types of SOLIDserver belong to the fourth generation hardware appliances, to the exception of the SDS-260, the hardware front panel LCD screen can be used in two ways:

1. To set up the iDRAC IP address.

Once the appliance has booted and is up and running, you have 30 seconds to press any arrow button to get into the edition mode. Through this mode you can edit the IP address, netmask and gateway. By default, it is configured with the IP address `192.168.0.120/255.255.255.0` and the gateway `192.168.0.0`.

2. To visualize the iDRAC IP address and DNS and DHCP basic configuration.

The LCD screen allows you to view at all times the hardware settings: the iDRAC IP (IPv4, IPv6 and DNS configuration), the error history, the MAC address, the host, model and user string name, the electric power used and the temperature.

This hardware generation of SOLIDserver comes embedded with a sub-component designed to remotely manage the appliance. This sub-component is called an iDRAC controller, or Integrated

Remote Access Controller. The iDRAC provides a dedicated management platform from where you can monitor the hardware appliance itself and configure SOLIDserver software.

For these reasons, the installation of these hardware appliance models requires to:

1. Configure an IP address for the iDRAC.
2. Connect to the iDRAC to configure SOLIDserver access IP address.
3. Once the software installation is over, you need to connect to SOLIDserver, add your license and set up the Internal module setup. For more details, refer to the chapter [Using SOLIDserver for the First Time](#).

Prerequisites

- SOLIDserver hardware appliance SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 or SDS-Blast series (SDS-4000, SDS-5000 or SDS-5500).
- A browser with the latest version of Java JRE installed.

1. Setting up the iDRAC IP Address

To configure a SOLIDserver software appliance on a hardware appliance SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 or SDS-Blast series (SDS-4000, SDS-5000 or SDS-5500) you need to set up an IP address to connect to the hardware appliance iDRAC. Once the iDRAC IP address is set, you can access it and configure SOLIDserver IP address.

To set up the iDRAC IP address you must:

1. Make sure you meet the prerequisites.
2. Choose a configuration method:
 - Either without LCD screen, using a keyboard and monitor. This is **mandatory for SDS-260 appliances**. For more details, refer to the section [Setting up the iDRAC IP Address Without LCD Screen](#).
 - Or from the LCD screen. This method can be used on **appliances SDS-550 and higher**. For more details, refer to the section [Setting up the iDRAC IP Address From the LCD Screen](#).

Prerequisites

- Plug in the appliance power supply cable(s):
 - For appliances SDS-260, you only need to plug in the power cable.
 - For appliances SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast series (SDS-4000, SDS-5000 and SDS-5500), you must plug in both power cables.
- Prepare at least two Ethernet cables:
 - One is dedicated to the iDRAC configuration.
 - The other one is used to set up the appliance physical interfaces. You can set up an Ethernet port failover with up to 4 cables.

Setting up the iDRAC IP Address Without LCD Screen

For hardware appliances SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 or SDS-Blast series (SDS-4000, SDS-5000 or SDS-5500), you can configure an access IP address for the iDRAC using a keyboard and monitor. This is the **only method for SDS-260 appliances**. Once

the iDRAC IP address is set, you can access it and configure the IP address of your SOLIDserver appliance.

The iDRAC is already configured with a default IP address - *192.168.0.120*, netmask - *255.255.255.0* and gateway - *192.168.0.0*. If this IP address is free and accessible on your network, you do not need to follow the procedure below, you can go straight to the section [Configuring SOLIDserver from the iDRAC](#).

To set up the iDRAC IP address without LCD screen

1. Plug in the Ethernet cables at the back of the appliance:
 - a. Plug in one Ethernet cable in the iDRAC plug, on the left-end side of the back panel.
 - b. Plug in up to four physical interface Ethernet cables.
2. Connect a monitor and keyboard to the appliance.
3. Press the button  to turn on the hardware appliance.
4. When the **BIOS splash screen** appears, hit **[F2]** on your keyboard to enter the system setup.

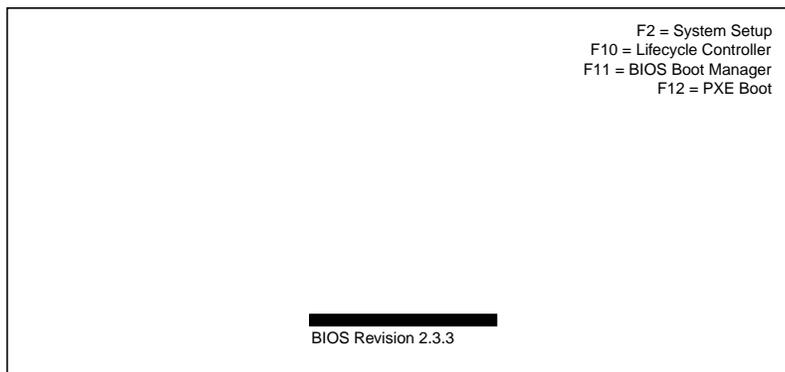


Figure 2.1. BIOS splash screen

The window **System Setup** eventually opens.

5. In the **System Setup Main Menu**, use the **down arrow** to highlight **iDRAC Settings** and hit **Enter**. The page **iDRAC Settings** opens.
6. Use the **down arrow** to highlight **Network** and hit **Enter**. The page **iDRAC Settings - Network** opens.
7. Use the **down arrow** to get to the section **IPV4 SETTINGS** and, if need be, edit the following fields:
 - a. In the field **Static IP Address**, set the IP address of your choice. You need it to connect to the iDRAC. The default IP address is *192.168.0.120*.
 - b. In the field **Static Gateway**, set the gateway of your choice. The default gateway is *192.168.0.0*.
 - c. In the field **Static Subnet Mask**, set the netmask of your choice. The default netmask is *255.255.255.0*.
8. When you have configured the three fields, hit the **tabulation** key to select the button **Back** in the window bottom right corner and hit **Enter** to get back to the page **iDRAC Settings**.

9. Hit **tabulation** to select the button **Finish** and hit **Enter**. The window **Warning / Saving Changes / Settings have changed. Do you want to save the changed?** opens.
10. Select **Yes** and hit **Enter**. The window **Success / Saving Changes / Settings saved successfully** opens.
11. Hit **Enter**. The menu **System Setup** opens.
12. Hit **tabulation** to highlight the button **Finish** and hit **Enter**. The window **Warning / Confirm Exit / Are you sure you want to exit and reboot?** opens.
13. Select **Yes** and hit **Enter**. The menu closes and the hardware appliance reboots.

Now you can connect to the iDRAC to configure SOLIDserver as detailed in the section [Configuring SOLIDserver from the iDRAC](#).

Setting up the iDRAC IP Address From the LCD Screen

For hardware appliances SDS-550, SDS-1100, SDS-2200, SDS-3300 or SDS-Blast series (SDS-4000, SDS-5000 or SDS-5500), you can set up an access IP address for the iDRAC from the LCD screen. You **cannot use this method for SDS-260 appliances**. Once the iDRAC IP address is set, you can access it and configure the IP address of your SOLIDserver appliance.

If you already configured the iDRAC IP address using a keyboard and monitor, following the section *Configuring the iDRAC IP address Without LCD Screen*, you can go straight to the section [Configuring SOLIDserver from the iDRAC](#).

The iDRAC is already configured with a default IP address - *192.168.0.120*, netmask - *255.255.255.0* and gateway - *192.168.0.0*. If this IP address is free and accessible on your network, you do not need to follow the procedure below, you can go straight to the section [Configuring SOLIDserver from the iDRAC](#).

To set up the iDRAC IP address from the LCD screen

1. Plug in the Ethernet cables at the back of the appliance:
 - a. Plug in one Ethernet cable in the iDRAC plug, on the left-end side of the back panel.
 - b. Plug in up to four physical interface Ethernet cables.
2. Press the button  to turn on the hardware appliance. During the boot sequence, the LCD screen displays *System booting*. The message disappears once the appliance has booted.
3. Press  to display the menu.
4. Press  to highlight *Setup*. Press  to enter the Setup menu.
5. *iDRAC* is highlighted, press . The message *Security Notice: Default network credentials in effect* appears.
6. Press  to enter the menu. Press  twice to highlight *Static IP*.
7. Press  to enter the menu. The default IP address is displayed: *IP: 192.168.0.120*. Set the IP address of your choice:
 - Use the arrows  and  to put the cursor over the digit you want to edit.
 - Once the cursor is on the digit you want to edit, press . The cursor blinks.
 - Once the cursor blinks, use the arrows to change the digit value and press  to commit the new value.

- Repeat these actions for every digit you want to change.

Once the full IP address matches your needs, press **▶** until the cursor highlights **>>** at the end of the IP address.

8. Press **✔** to commit the IP address. The subnet mask is displayed: *Sub: 255.255.255.000* . Set the netmask of your choice using the technique detailed in step 7.

Once the full netmask matches your needs, press **▶** until the cursor highlights **>>** at the end of the subnet mask.

9. Press **✔** to commit the subnet mask. The gateway is displayed: *Gtw: 192.168.000.000* . Set the gateway of your choice using the technique detailed in step 7.

Once the gateway matches your needs, press **▶** until the cursor highlights **>>** at the end of the subnet mask.

10. Press **✔** to commit the gateway IP address. The DNS configuration menu appears: *Setup DNS: Yes*.
11. Press **▶** to highlight *No* and press **✔** to skip the DNS configuration. The screen displays: *Save: Yes / No*, *Yes* is highlighted.
12. Press **✔** to save your configuration. The LCD screen is now empty.

Now you can connect to the iDRAC to configure SOLIDserver as detailed in the section [Configuring SOLIDserver from the iDRAC](#) below.

2. Configuring SOLIDserver from the iDRAC

On SOLIDserver hardware appliances SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast series (SDS-4000, SDS-5000 and SDS-5500) you must configure SOLIDserver software:

1. Connect to the iDRAC platform, following the procedure [To connect to SOLIDserver from the iDRAC platform](#).
2. Connect to SOLIDserver using the console available in the iDRAC GUI and configure SOLIDserver IP address, following the procedure [To configure SOLIDserver from the iDRAC platform](#).

To connect to SOLIDserver from the iDRAC platform

1. Open any browser that has Java installed, in the URL field type in **https://<iDRAC-configured-IP-address>**.
2. Accept the certificate. The iDRAC login page opens.
3. In the field **Username**, type in *root*.
4. In the field **Password**, type in *calvin*.
5. Hit **Enter**. The iDRAC homepage opens.
6. Launch the **Virtual console**.

If the iDRAC is in version 8:

- a. Tick the box **Keep Default Password** and click on **Continue**. The iDRAC homepage opens: **System Summary**. It details the hardware information.

- b. In the panel **Virtual Console Preview**, click on **Launch**.
 - c. Accept to save the file *viewer.jnlp* to be able to launch the console.
 - d. Open the file *viewer.jnlp*¹.
 - e. In the pop-up window *SecurityWarning*, click on **Continue**.
 - f. In the pop-up window *Do you want to run this application?*, click on **Run**.
 - g. In the pop-up window *Warning-Security*, click on **Run** to accept the certificate. The console opens.
7. When SOLIDserver installation is complete, the page **WELCOMETO SOLIDSERVER** appears.

```
#####
#####
                                WELCOME TO SOLIDSERVER
#####
#####

.
Starting inetd.
Fri Jan 16 17:04:10 UTC 2015

SOLIDserver access

login:
```

Figure 2.2. SOLIDserver login page

If an error message is displayed on the screen, you need to reimage SOLIDserver. For more details, refer to the guide *SOLIDserver_Hardware_4th-Gen_Reimaging-.x.x.x.pdf* available on our website ².

To configure SOLIDserver from the iDRAC platform

1. From the page **WELCOMETO SOLIDSERVER**, log in using the login *root* and no password. The configuration window opens.
2. In the **Main Menu**, the line **S Quick Start** is highlighted.
3. Hit **Enter**. The page **Quick Start - Configure basic network settings** opens.
4. Hit **Enter** to confirm the quick start configuration. The menu **Physical interfaces** opens.
5. In the menu **Physical Interfaces**, select one or several interfaces. *[*]* indicates which interface(s) is selected. Use the keyboard arrows or the digits to highlight the interface of your choice and hit the **space** key to select or deselect an interface.

If you have plugged in several physical interfaces Ethernet cables, you are configuring an Ethernet port failover.

The interfaces are named and numbered: *bgeX*, *igbX* and/or *ixX*³, where X is a digit. Their status is displayed between brackets: (*active*) or (*no carrier*).

6. Hit **Enter** to confirm the interface(s) selection. The menu **Network configuration** opens.

¹If the file *.jnlp* has not been automatically opened by Java, you have to associate it manually with the file *javaws.exe* located in the folder *jre\bin* of the appropriate Java version.

²At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

³*ixX* interfaces correspond to 10 Gb ports.

7. Configure the access to SOLIDserver:
 - a. Configure the interface IP address on the line **IP:** . Hit the **down arrow** to get to the next line.
 - b. Configure the interface netmask on the line **Netmask:** . Hit the **down arrow** to get to the next line.
 - c. Configure the interface default gateway on the line **Gateway:** . Hit the **tabulation** key to highlight OK.
8. Hit **Enter** to confirm your configuration and go back to the **Main Menu**.
9. The line **C Commit modifications to system** is highlighted.
10. Hit **Enter** to save the whole configuration. The window **Do you really want to commit modifications to system?** opens.
11. The button **Yes** is highlighted.
12. Hit **Enter** to commit your choice. When the configuration is saved, the window **Configuration applied** opens.
13. Hit **Enter** to close the window. The **Main Menu** is visible again.
14. Hit **E** to select **EXIT** and hit **Enter**. The configuration window closes.

Now your configuration is complete and you can connect to SOLIDserver software using the IP address you configured.

SOLIDserver Software First Installation from a USB Flash Drive

If you purchased a SOLIDserver software appliance, the software image can be saved on a USB flash drive that allows to install it on your own hardware appliance, a non-Efficient IP appliance.

To be able to install SOLIDserver, your hardware appliance must match the following requirements:

Table 2.1. SOLIDserver software installation requirements on your own hardware

	SDS-50	SDS-260	SDS-550	SDS-1100	SDS-2200	SDS-3300
VCPU	1	2	2	4	4	8
RAM	2 GB	4 GB	8 GB	8 GB	16 GB	32 GB
Virtual disk size	> 32 GB	> 32 GB	> 32 GB	> 64 GB	> 64 GB	> 64 GB
RAID Controller	Any card compatible with FreeBSD version 11.					
NIC	Any NIC, except if your license includes Guardian; in that case you need an Intel Ethernet i350.					

To install SOLIDserver on a hardware appliance using a USB flash drive and SSH

1. Make sure your hardware appliance meets the [requirements](#).
2. Prepare a bootable USB flash drive with SOLIDserver image on it.
3. Plug the USB flash drive in the hardware appliance.
4. Configure the server booting process to make your USB flash drive the first booting device.
5. Install SOLIDserver on the appliance:
 - a. Select **USB KEY (<your USB flash drive>) INSTALLATION on <your appliance>**.

- b. Hit **Enter** to confirm.
 - c. Select the SOLIDserver image you want to install.
 - d. Hit **Enter** to confirm.
 - e. Select **CONFIRM THE INSTALLATION**.
 - f. Hit **Enter** to confirm. A progress bar appears to display the installation process.
 - g. Once the installation is completed, hit **Enter** to reboot the appliance.
6. Put the server HDD back as **1st Drive** and **1st Boot Device**.

Basic Network Configuration

After installing SOLIDserver via the menu *Quick Start* detailed above, you can:

- Use a terminal to complete the configuration (set a hostname and default gateways, name physical interfaces...), as detailed in the section [Setting the Basic Network Configuration Using a Terminal](#).
- Reset the network settings from the LCD screen, as detailed in the section [Resetting the Default Network Settings From the LCD Screen](#).

SOLIDserver software and hardware appliances can be configured using a terminal accessible through Command-Line Interface (CLI). Any appliance provides it:

- Software appliances provide CLI thanks to:
 - SOLIDserver IP address that you can use to open a shell session via a terminal emulator.
 - A VMware console of you installed SOLIDserver on VMware.
- Hardware appliances provide CLI in different ways, it depends on the model:
 - For appliances SDS-50, SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000, you can:

Plug the hardware serial port to your computer and open a terminal emulator that provides serial connection sessions.

- For appliances SDS-260, SDS-550, SDS-1100, SDS-2200, SDS-3300 and SDS-Blast series (SDS-4000, SDS-5000 and SDS-5500), you can:

Connect to the iDRAC GUI and open the virtual console from the page Overview/Virtual console

- For any appliance, except the SDS-50 you can:

Plug a monitor and keyboard into the hardware appliance VGA and USB ports.

Setting the Basic Network Configuration Using a Terminal

By default there is already an interface configured for SOLIDserver image, it has the IP address **192.168.1.1** and the netmask **255.255.255.0**. You can change both according to your need when configuring the basic network settings or during the first installation.

In the procedure below, we configure a virtual interface, a physical interface, an IP address, a netmask, a gateway and the IP address of the DNS resolver.

To configure basic network setting through a terminal

1. Connect to SOLIDserver appliance [through CLI](#).
2. Once you are connected, the page **WELCOME TO SOLIDserver** opens.

```
#####  
#####  
  
WELCOME TO SOLIDSERVER  
  
#####  
#####  
  
.  
Starting inetd.  
Fri Jan 16 17:04:10 UTC 2015  
  
SOLIDserver access  
  
login:
```

Figure 2.3. SOLIDserver login page

3. Log in using the login *root* and no password. The **Main menu** appears.
4. The line **N Network Configuration** is highlighted.

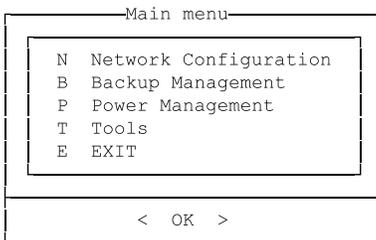


Figure 2.4. Main menu

Hit **Enter**, the menu **Network configuration** opens.

5. The line **V Virtual interfaces** is highlighted.

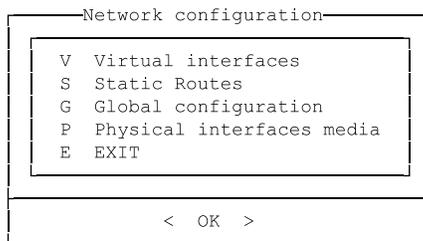


Figure 2.5. Network configuration

Hit **Enter**, the menu **Virtual interfaces** opens.

6. The line **1 DEFAULT_INTERFACE** is highlighted.

First Installation and Basic Configuration

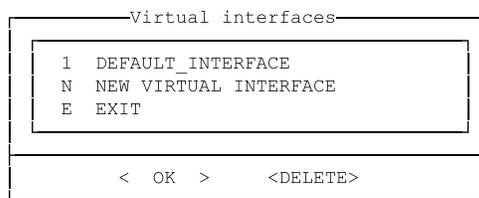


Figure 2.6. Virtual interfaces

Hit **Enter**, the menu **Virtual interface name** opens.

7. Edit the interface **Name** if you want.

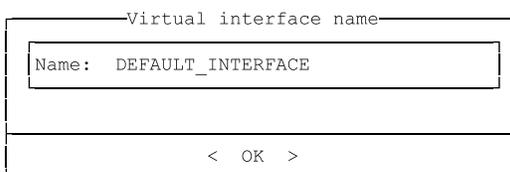


Figure 2.7. Virtual interface name

Hit **Enter**, the menu **Physical interfaces** opens.

8. In the menu **Physical interfaces**, select one or several interfaces.

[*] indicates which interface(s) is selected. Use the keyboard arrows or the digits to highlight the interface of your choice and hit the **space** key to select or deselect an interface.

If you have plugged in several physical interfaces Ethernet cables, you are configuring an Ethernet port failover. The interfaces are named and numbered. Their status is displayed between brackets: (*active*) or (*no carrier*).

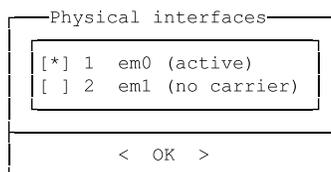


Figure 2.8. Physical interfaces⁴

Hit **Enter** to confirm the interface(s) selection, the menu **IP addresses list** opens.

9. The line **1 192.168.1.1 255.255.255.0** is highlighted. The IP address *192.168.1.1* and the netmask *255.255.255.0* are the default values.

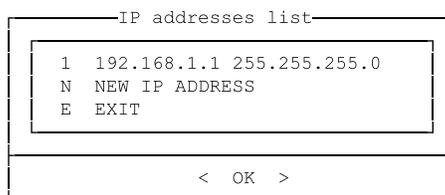


Figure 2.9. IP addresses list

⁴The interface name depends on the appliance. In this example, the interfaces name can be found on appliances SDS-250, SDS-500, SDS-1000, SDS-2000 and SDS-3000.

Hit **Enter**, the menu **IP addresses configuration** opens.

10. Edit the **Address** and **Prefix/Netmask** fields according to your needs.

```
IP addresses configuration
-----
Address:          192.168.1.1
Prefix/NetMask:  255.255.255.0
Specific Route:
802.1q tag:
-----
< OK >
```

Figure 2.10. IP addresses configuration

Hit **Enter** to save your changes, the menu **IP addresses list** menu opens again.

11. Hit the key **E** to select **E EXIT** and it **Enter**. The menu **Network configuration** opens.
12. Hit **G** to select **G Global configuration**.

```
Network configuration
-----
V Virtual interfaces
S Static Routes
G Global configuration
P Physical interfaces media
E EXIT
-----
< OK >
```

Figure 2.11. Network configuration

Hit **Enter**, the menu **Global configuration** opens.

13. Edit the **Hostname** if need be. Edit the **Default IPv4 gateway** and **1st DNS resolver** according to your needs.

```
Global configuration
-----
Hostname:          nsl.test.corp
Default IPv4 gateway: 192.168.1.254
Default IPv6 gateway:
1st DNS resolver:  192.168.1.254
2nd DNS resolver:
-----
< OK >
```

Figure 2.12. Global configuration

Hit **Enter** to save your changes, the menu **Network configuration** opens again.

14. Hit **E** to select **E EXIT** and it **Enter**. The menu **Main menu** opens.
15. The line **C Commit modifications to system** is highlighted.

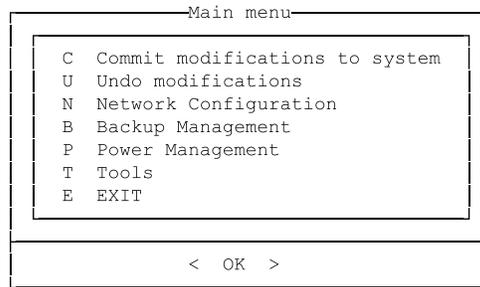


Figure 2.13. Main menu before saving your configuration

Hit **Enter** to save the whole configuration.

16. In the confirmation window, the button **Yes** is highlighted.



Figure 2.14. Confirmation window

Hit **Enter** to commit your configuration.

17. When the configuration is saved, the window **Configuration applied** opens.

Hit **Enter** to close the window. The **Main Menu** is visible again.

18. Hit **E** to select **E EXIT** and it **Enter**. The configuration window closes, the page **SOLIDserver access** is visible again.

Now your configuration is complete and you can access your SOLIDserver through the browser of your choice. Make sure the browser version complies with the [prerequisites](#) mentioned above.

Resetting the Default Network Settings From the LCD Screen

If you purchased the hardware appliance SDS-250, SDS-500, SDS-1000, SDS-2000 or SDS-3000 you can reset the system to its original network configuration parameters from the appliance LCD screen.

The default configuration of the third generation of SOLIDserver hardware appliances is the following:

- IP address/Netmask: `192.168.1.1 / 255.255.255.0` .
- Gateway: `0.0.0.0` .
- Physical interface: `eth1`.

Keep in mind that **resetting the hardware appliance default configuration automatically reboots it**.

To reset the basic network configuration via LCD

1. If the SOLIDserver is not turned on, push **⏻**. During the boot sequence, the LCD screen displays *Device booting*. Once SOLIDserver is started, a timer is displayed. Press **⏻** to enter in the setup.
2. Push **⏻** until the **EXIT** menu appears. Then press **⏻** to enter in the **RESET** menu. The message *WARNING: net conf will be lost* is displayed.
3. To commit the reset, push **⏻**. The message *Restarting appliance...* appears. If you do not want reset, push **⏻** to discard.

Chapter 3. Using SOLIDserver for the First Time

When using SOLIDserver for the first time, you need to:

1. Log in.
2. Request and add a license.
3. Define the modules interaction, or internal module setup.
4. Configure SOLIDserver on your network.
5. Create groups of users, add users and set their rights.

Connecting to SOLIDserver

No matter the browser you choose to use, to access SOLIDserver you need to follow the procedure below.

If you defined a hostname for your SOLIDserver in your DNS, you can use its name in the URL field rather than the IP address.

To connect to SOLIDserver for the first time

1. Open your browser, in the URL field type in **https://<SOLIDserver-configured-IP-address>**.
2. Hit **Enter**. The browser displays a security warning because the default certificate is in use.

Your browser probably identified that the certificate is not from a trusted certifying authority and that the hostname on the certificate is invalid or does not match the name of the site.

3. Accept the certificate. SOLIDserver login page appears:



Figure 3.1. First connection to SOLIDserver

4. In the field **Login**, type in *ipmadmin*. The default superuser login.
5. In the field **Password**, type in *admin*. The default superuser password.
6. Click on **OK** to complete the operation. The page **Main Dashboard** opens, SOLIDserver homepage.

On the dashboard the gadget **System Information** indicates that there is *No license installed*, you must request a license and add it to manage SOLIDserver.

Requesting and Adding a License

During the first connection, you need to request a license from EfficientIP. They will generate and send you a valid license key that you must add to the GUI to activate the license and use SOLIDserver.

Each license key is unique and specific to one SOLIDserver appliance, you cannot use the same license key on several appliances.

To request a license key

1. Retrieve the request license key.
 - a. On SOLIDserver **Main Dashboard**, a set of gadgets is visible.
 - b. In the gadget **System Information**, click on the link **Request license**. The wizard **Request license** opens.
 - c. Read the **Software License Agreement** and click on **NEXT**. The next page opens.
 - d. Copy the content of the field **Key**, you need it to fill out the request license form.
 - e. Click on **OK** to close the wizard.
2. Send the request key to Efficient IP.
 - a. Go to the page <http://www.efficientip.com/license-request/> and fill out the *Request Your License* form.
 - b. In the fields **FIRST NAME**, **LAST NAME**, **EMAIL**, **COMPANY**, **PHONE NUMBER** and **COUNTRY NAME**, specify your contact details. All these fields are required.
 - c. In the field **LICENSE PERIOD REQUEST**, specify the desired license length: *1 month*, *2 months*, *3 months*, *6 months* or *Permanent*. This field is required.
 - d. If you selected *Permanent*, you must fill in the field **CONTRACT NUMBER (IF PERMANENT LICENSE)**.
 - e. In the field **REQUEST KEY**, paste your request key or the content of your request key file. This field is required.
 - f. In the field **NUMBER OF EXTERNAL MANAGED SERVERS (MVSM, IF ANY)**, specify the total number of servers - DNS/DHCP/... - you intend to manage from SOLIDserver.
 - g. In the section **OPTIONAL MODULE**, tick all the optional modules you might need: *NETCHANGE*¹, *DEVICE MANAGER*, *SPX* or *DNS GUARDIAN*.
 - h. If relevant, fill in the field **IF REQUESTER IS OTHER THAN THE END CUSTOMER, PLEASE PROVIDE YOUR CONTACT INFORMATION (NAME, COMPANY, EMAIL, PHONE)**.
 - i. Click on **SUBMIT** to send us your information.

¹If you do not tick this box, you are using NetChange basic options, or NetChange-IPL.

Once EfficientIP has answered your request and sent you a license key, you can add it to the appliance to activate your licence as detailed below. Note that **you cannot activate the license if the appliance is not on time.**

To activate a license

1. From the EfficientIP email response to your license request, copy the license key.
2. Connect to SOLIDserver using the superuser credentials. The page **Main Dashboard** opens.
3. In the gadget **System Information**, click on the link **Add license**. The wizard opens.
4. Read the **License Agreement** and click on **NEXT**. The page **Add a license** opens.
5. In the field **License**, paste the license key.
6. Click on **OK** to complete the operation. The page refreshes, in the gadget **System Information** the *License type* is now displayed and all the modules that come with your license are visible.

If you need to renew your license, because you want more services or the temporary license is about to expire, refer to the chapter [Managing the Licenses](#).

Defining the Internal Module Setup

Once the license is installed, and before making any further network configurations, we recommend that you set the *Internal module setup*.

This setup allows you to enable the interaction between the IPAM, DNS and DHCP modules. That way you can manage your resources and objects on one page and update them in other modules.

To configure the internal module setup from the configuration gadget

1. In the sidebar, go to  **Dashboard** > **Main Dashboard**, SOLIDserver homepage.
2. At the bottom of the gadget **SOLIDserver Configuration Checklist**, next to *Internal module setup*, click on **Configuration**. The wizard **Internal module setup** opens.
3. If you want to activate all the module interactions:
 - a. In the drop-down list **Architecture**, select *IPAM*.
 - b. Tick the box **Use DNS**.
 - c. Tick the box **Use DHCP**.
4. If you only want to activate the DNS, in the drop-down list **Architecture**, select *DNS only*.
5. Click on **OK** to complete the operation. The wizard closes. The Main Dashboard is visible again, the *Internal module setup* is marked .

At any point, you can edit the internal module setup from the module Administration.

To configure the internal module setup from the module Administration

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Internal module setup**. The wizard opens.
3. If you want to activate all the module interactions:

- a. In the drop-down list **Architecture**, select *IPAM*.
 - b. Tick the box **Use DNS**.
 - c. Tick the box **Use DHCP**.
4. If you only want to activate the DNS, in the drop-down list **Architecture**, select *DNS only*.
 5. Click on to complete the operation. The wizard closes.

Configuring SOLIDserver Network and Services

Before using SOLIDserver you need to complete its configuration. The appliance *Main Dashboard* allows to quickly set network and service configurations:

1. The gadget **SOLIDserver Configuration Checklist** provides an overview of key configurations.

What is marked  has not been configured yet. You can click on each link to configure everything that suits your needs. For more details, refer to the section [SOLIDserver Configuration Checklist](#).

2. The gadget **General Information** provides shortcuts toward services and appliance settings.

- **Services** lists key services.  services are not configured yet,  are not started yet and  are configured and running. Clicking on configured services allows to start/stop them, clicking on services not configured yet opens the page *Services configuration* where you can configure them. For more details, refer to the chapter [Configuring the Services](#).

- **Hostname, IP Addresses, Default gateways** and **Status** return appliance basic network configuration details. Clicking on any value opens the page *Network Configuration* where you can edit or set them. For more details, refer to the chapter [Configuring the Network](#).

- **SOLIDserver role** is the appliance role on the page *Centralized management*. For more details, refer to the chapter [Centralized Management](#).

3. Make sure that you configured and started the service **NTP**, it is crucial for DHCP, DNS and High Availability management. For more details, refer to the chapter [Configuring the Time and Date](#).

Creating SOLIDserver Users

Being logged as *ipmadmin*, the superuser, you probably need to add users and set their privileges. Within SOLIDserver, users rights and resources depend on the group they belong to.

The most privileged group is *admin*, to which *ipmadmin* belongs. The other groups, even with all the permissions granted, are not able to perform some operations. If an operation can only be performed by a user from the group *admin*, it is specified in the procedure.

To configure specific users permissions you need to:

1. Create a group of users and configure it with the permissions that suit your needs.
2. Add one or more users to that group.
3. Set authentication rules if these users are not local users.

Once users are created and configured, do not forget to grant them access to existing resources if need be. For more details regarding users and groups, refer to the part [Rights Management](#).

Chapter 4. Understanding the GUI

SOLIDserver centralizes all the operations in a unified Graphical User Interface (GUI) divided into modules with common management principles and navigation logic.

All management pages share common elements, detailed in the image below.

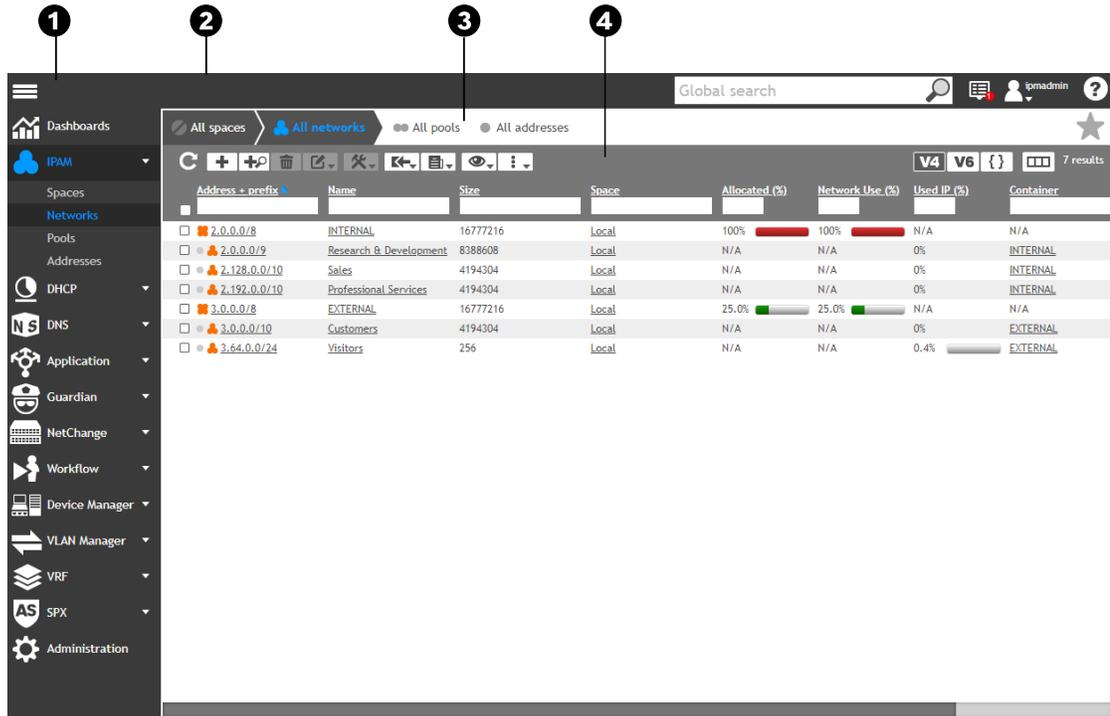


Figure 4.1. Common elements on listing pages in the GUI

- 1 The sidebar gives access to the SOLIDserver dashboards, the modules, the Tree view and the Smart folders. For more details, refer to the section [Sidebar](#).
- 2 The top bar gives access the Global search, your settings through the menu *My account*, all *Notifications* and the menu *Help*. For more details, refer to the section [Top Bar](#).
- 3 The breadcrumb is a navigation bar available on all pages of all modules except *Administration*, where it is only displayed on some pages. It provides direct and hierarchical access to the module objects. For more details, refer to the section [Breadcrumb](#).
- 4 The menu is displayed on every page. Its content differs on each page. For more details, refer to the section [Menu](#).

SOLIDserver Main Dashboard

Once you are connected, SOLIDserver *Main Dashboard*, the appliance home page, opens. It contains the welcome banner and the *Main Dashboard* right under it. You can access this page from anywhere in the appliance if you click on the button  **Dashboards** in the sidebar.

The *Main Dashboard* contains one or several gadgets depending on the user connected. For more details, refer to the section [Gadgets Displayed by Default](#).

You can edit the *Main Dashboard* welcome banner with a different message or even an image. For more details, refer to the chapter [Customizing the GUI](#).

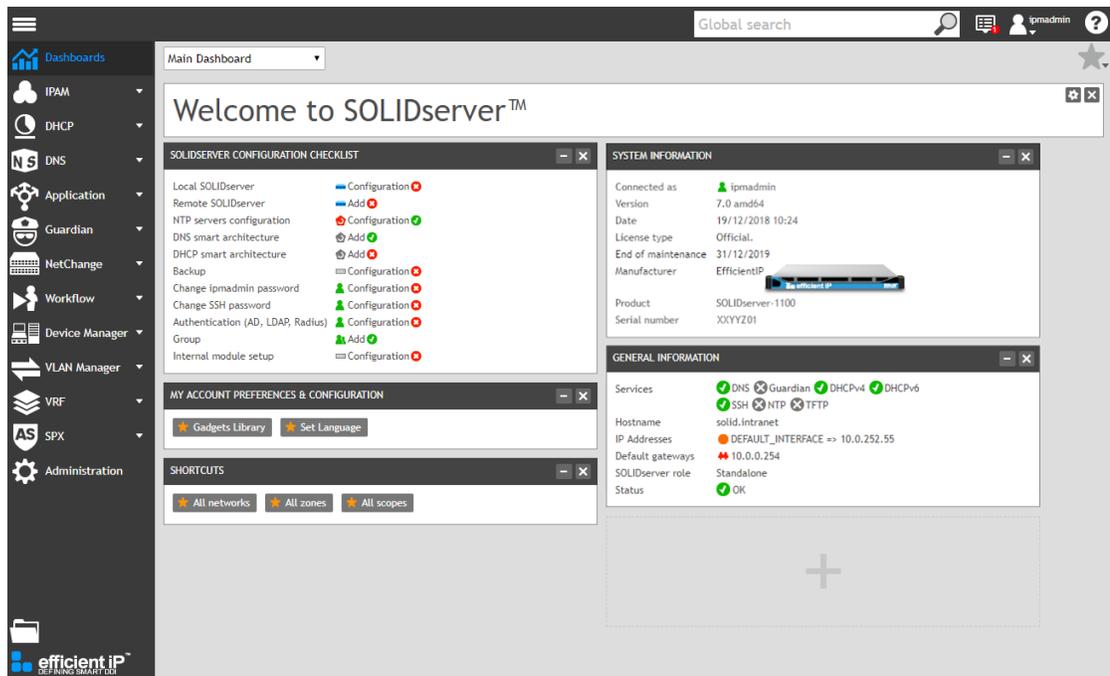


Figure 4.2. The Main Dashboard of the superuser session

In the module **Dashboards**, you can access the dashboard of a specific module through the drop-down list on the page *Main Dashboard*.

On the dashboards, you can add the gadgets of your choice. For more details, refer to the chapter [Managing Gadgets](#).

Sidebar

The sidebar allows to access all the modules as well as the *Tree view* and the *Smart folders*.

The toggle button  allows to either display:

- a collapsed menu with only the icons of each module. Hover over the icons to display the name of a module and the list of its pages.
- an expanded menu with the icon and name of each module. Below the name of a module, the list of its pages is displayed.

Modules

Each module has a dedicated section in the sidebar.

When you navigate from one module to another, the latest page visited is saved and you can open it again if you go back to the module.

Clicking on the icon or name of the module you are currently in opens the page of the top level.

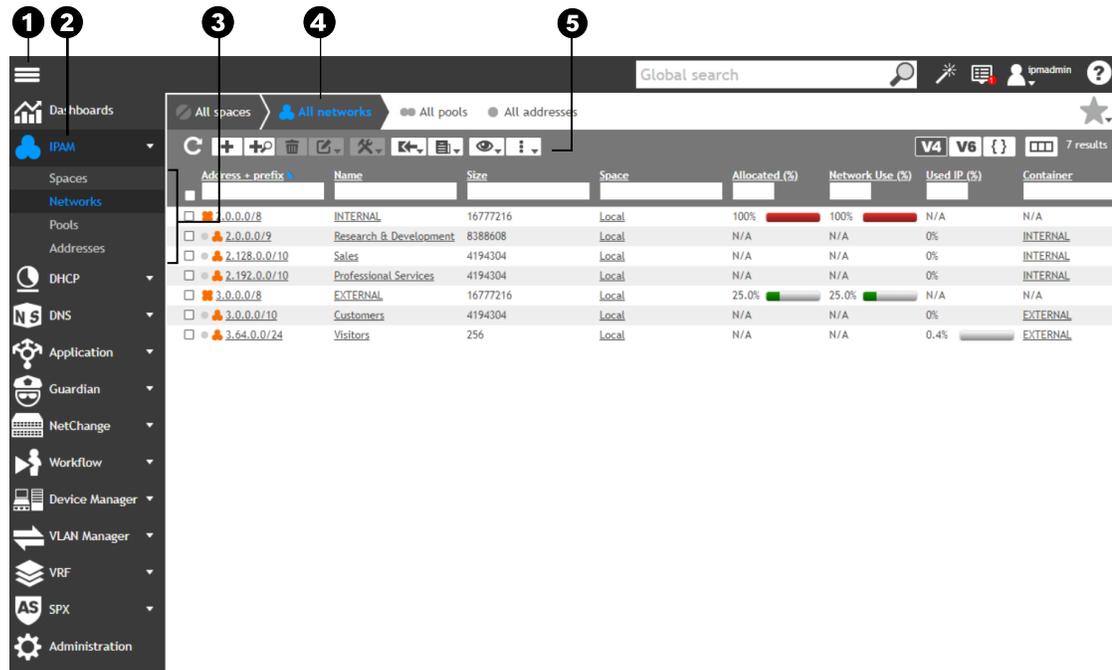


Figure 4.3. Overview of the page All Networks in the module IPAM

- 1 When the sidebar is expanded, each module icon and name are displayed.
- 2 The button highlighted in blue indicates the module you are currently in.
- 3 The module you are currently in is displayed. The current page is highlighted in blue.
- 4 In the breadcrumb, the current page is also highlighted in blue.
- 5 The menu contains some options specific to the module. For more details, refer to the section [Menu](#).

To open a specific page of a module when the sidebar is expanded

1. If the sidebar is expanded, next to the name of the module of your choice, click on the arrow. The list of the pages of this module is displayed.
2. Click on the name of the page you want to open. The page opens.

In the image above, to list the networks, in the sidebar, go to **IPAM > Networks**. The page *All networks* opens.

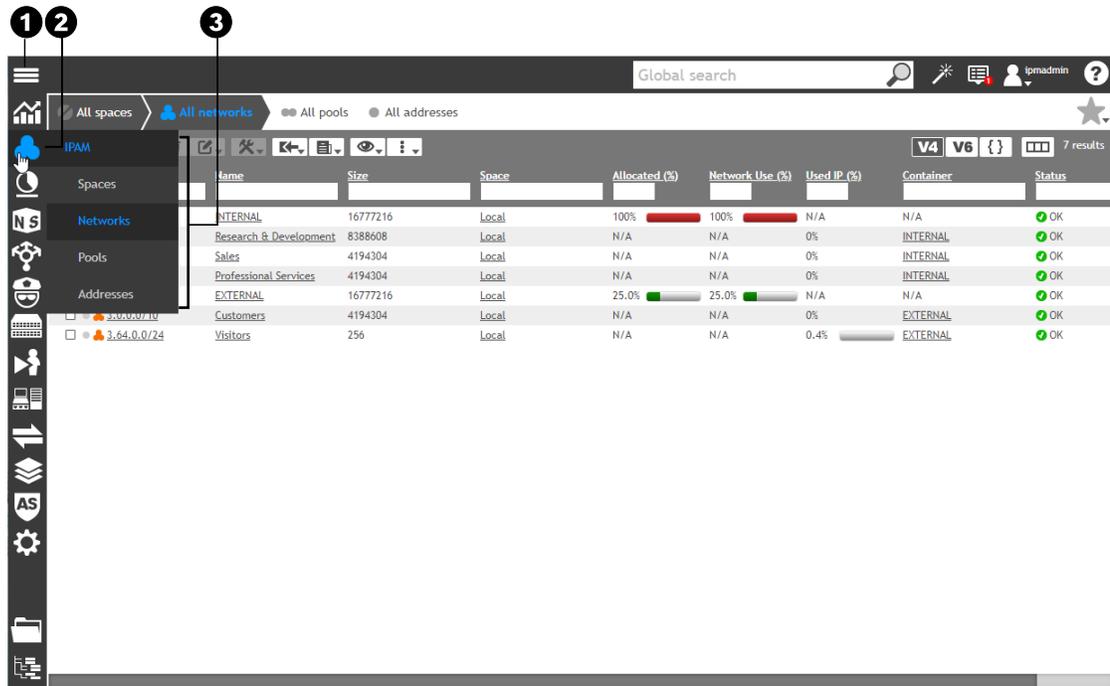


Figure 4.4. Overview of the page All Networks in the module IPAM

- 1 When the sidebar is collapsed, only the module icons are displayed.
- 2 The button highlighted in blue indicates the module you are currently in.
- 3 When you hover over the icon of a module, its name and the list of the pages it contains are displayed.

To open a specific page of a module when the sidebar is collapsed

1. Hover over the icon of the module of your choice. The name and the list of the pages of this module are displayed.
2. Click on the name of the page you want to open. The page opens.

In the image above, to list the networks, in the sidebar, go to **IPAM > Networks**. The page *All networks* opens.

Keep in mind that in addition to the menu entries in the sidebar, the breadcrumb allow to access the pages of the module. For more details, refer to the section [Breadcrumb](#).

The module *Administration* has a dedicated sidebar and a homepage, *Admin Home*.

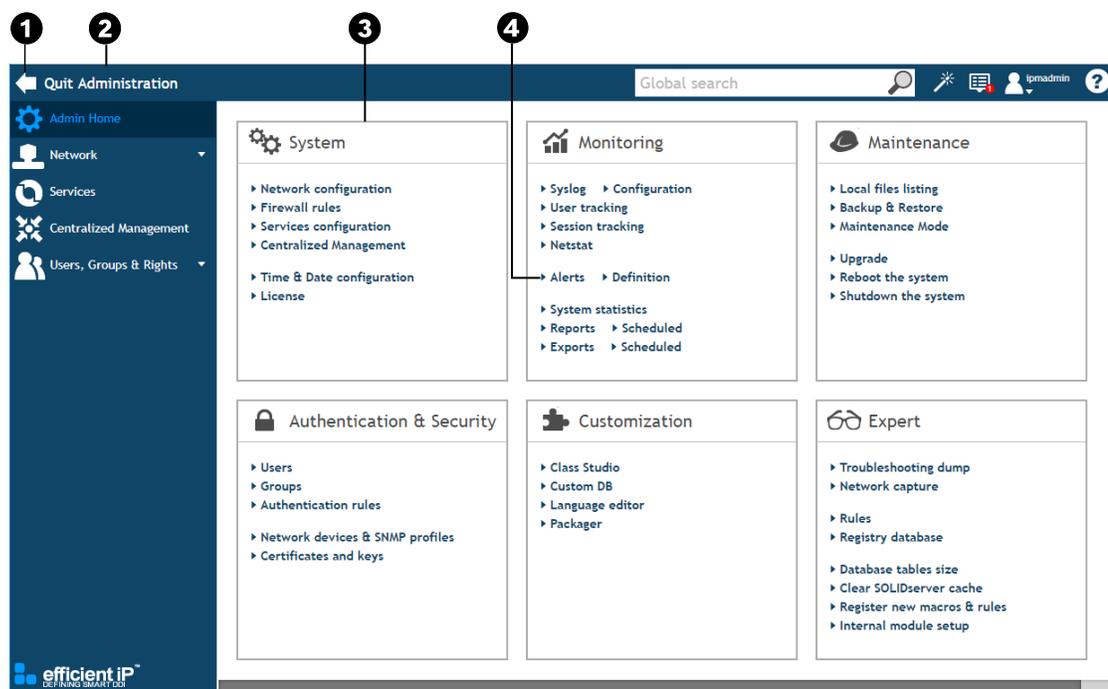


Figure 4.5. Overview of Admin Home, the homepage of the module Administration

- 1 The button *Quit Administration* allows to leave the module *Administration* and to go back to the other modules.
- 2 The sidebar cannot be reduced in the module *Administration*. It allows to navigate between specific pages of the module without displaying the homepage.
- 3 *Admin Home* is divided into six sections. In each section, you find links to open the different pages of the module.
- 4 When there are two links on one line, it means that the pages they open are linked to the same object. For example, for the links *Alerts* and *Definitions*, the link *Definition* opens the page *Alerts Definition*.

To open a specific page in the module Administration

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In a section, click on the link of your choice. The page opens.

Tree View

The Tree view is a side panel that allows to display and access the data of the modules *IPAM*, *DHCP* and *DNS*. When one of these modules is open, the button **Tree view** is available at the bottom of the sidebar and allows to display a panel that contains the following information:

DHCP

This section is divided among DHCP servers, DHCP scopes, DHCPv6 servers and DHCPv6 scopes. Open each type to display and access the servers and scopes you manage in the module DHCP.

DNS

This section is divided among servers and zones. Open each type to display and access the servers and zones you manage in the module DNS.

IPAM

This section is divided among the spaces you manage. Open each space sub-section to display and access the objects it contains in the module IPAM.

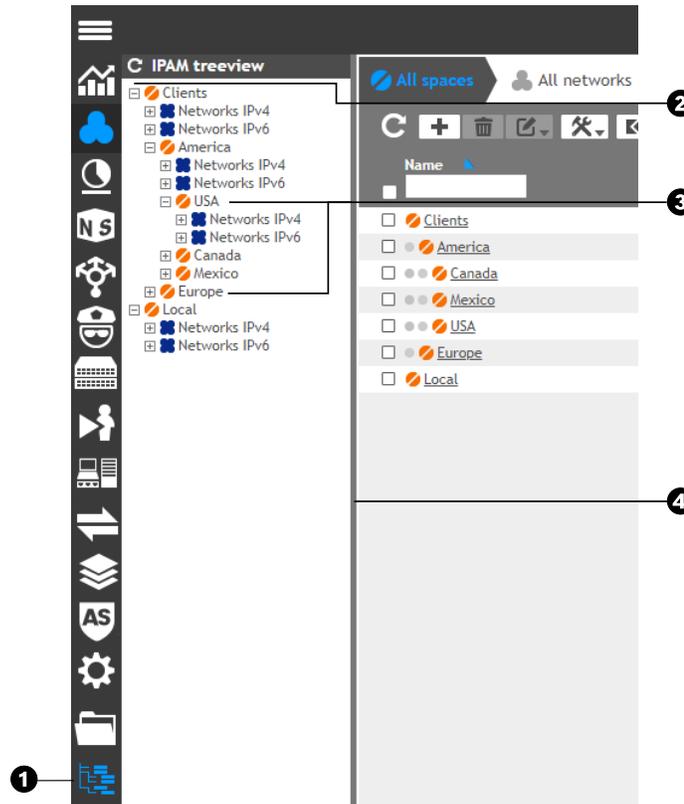


Figure 4.6. Overview of the IPAM tree view

- ❶ This button allows to open or close the Tree view. The button is highlighted in blue when the Tree view is open.
- ❷ This button allows to refresh a section.
- ❸ These icons indicate that a branch of the Tree view hierarchy is opened or closed . Click on it to display or hide the content.
- ❹ The Tree view panel can be dragged open according to your needs. Double-click on the right edge of the Tree view to close it.

To open the Tree view

1. In the sidebar, go to the module for which you want to display the Tree view: IPAM, DHCP or DNS .
2. In the sidebar, click on Tree view. The Tree view opens.
3. Each final line is a shortcut to the page listed. The page does not open if it is empty.

To refresh the content of the Tree view

1. Open the Tree view.
2. In the top left corner of the Tree view, click on **C**.

To expand or collapse the Tree view

1. Put your mouse over the right-end side of the Tree view. The pointer changes shape, a darker gray line appears.
2. To expand the panel, drag the line to the right.
3. To collapse the panel, double-click on the line or slide it the left. The Tree view collapses.

To close the Tree view

- In the sidebar, click on  **Tree view**. The Tree view closes.

If you open a module for which the tree view is not available, the panel closes automatically.

Smart Folders

This section contains all the smart folders you created, listed in the ASCII alphabetic order. It provides links following the hierarchy you set towards the objects you configured in these folders. The button  **My Smart Folders** is available at the bottom of the sidebar and allows to display a panel that contains your smart folders. For more details regarding smart folders refer to the chapter [Managing Smart Folders](#).



Figure 4.7. Overview of the Device Manager smart folders

- 1 This button allows to open or close the Smart Folders panel. The button is highlighted in blue when the Smart Folders panel is open.

- 2 These icons indicate that a branch of the Smart folders hierarchy is opened  or closed . Click on it to display or hide the content.
- 3 This button  allows to refresh a section.
- 4 The Smart Folders panel can be dragged open according to your needs. Double-click on the right edge of the panel to close it.

Top Bar

The top bar is available on all pages and gives access the *Global search*, your settings through the menu *My account*, the *Notifications* and the *Help*.



Figure 4.8. Overview of the Top bar

- 1  Global search, refer to the section [Global Search](#)
- 2  Notifications, refer to the section [Notifications](#)
- 3  My Account, refer to the section [My Account](#)
- 4  Help, refer to the section [?](#)

Global Search

SOLIDserver includes a search engine called Global search that allows you to **perform searches in the database of several modules at once**.

This Global search field is available from any page. You can look for data using full names, values or look for fragments of information (some letters of a name for instance) to find all the partial matches.



Figure 4.9. An example of search results performed using Global search

- 1 Type in the field the data you are looking for.

- 2 Click on the magnifying glass to perform the search.
- 3 All matching results are returned in sections dedicated to all the relevant modules and objects and preceded by their dedicated icon.
- 4 This button allows you to resize the window Global search.
- 5 To display more information in the window Global search, click on any result line. The data loads under the object line.
- 6 The button  allows to access the properties page of the object.
- 7 This button allows you to close the window. Once you performed a search, the window remains open above the page you are currently on, until you close it.

Note that Global search:

- Looks for matching results in the modules **IPAM, DHCP, DNS, NetChange, Device Manager** and two pages of **Administration**. It returns the following objects:
 - From the IPAM: IPv4 and IPv6 spaces, networks, pools and IP addresses, even in their compressed form.
 - From the DHCP: IPv4 and IPv6 smart servers, physical servers not managed via a smart architecture, scopes, ranges, statics, leases and ACLs.
 - From the DNS: smart servers, physical servers not managed via a smart, views, zones and resource records.
 - From NetChange: network devices, ports, device address and discovered items. To find discovered items, you have to search for their MAC or IP address. The result specifies if the discovered item was discovered via an *interco* port or not.
 - From Device Manager: devices, ports and interfaces.
 - From Administration: groups of users and users.
 - From all modules: Hostnames or MAC addresses.
- Returns results saved on names or values. For instance, you can find an IP address using its address, MAC address or name.
- Displays all the matching results following the order of the modules: **IPAM, DHCP, DNS, NetChange, Device Manager** and finally **Administration**.
- Displays all the matching results respecting the internal hierarchy of the module. So if you look for an IP address, Global search returns first the network(s) from level 0 to *n*, then the pool(s) and finally the IP address(es) matching your search.
- Does not return VLAN information, whether it is managed from NetChange or VLAN Manager.

To perform a global search

1. In the Top bar, in the field **Global search** type in the data you are looking for.
2. Click on  or hit Enter on your keyboard to perform the search.
3. The window **Global search** opens under the field and displays the results found in each module.
 - a. To display more information, click on any result line. <more information loads under the object line.
 - b. To access the object properties page in the module where it is managed, click on .

Notifications

SOLIDserver keeps track of all the operations performed during the last week in the dedicated window *Notifications*. Whether the operation was successful or not, this window lists the services run and allows to see and export the final report of the operation.

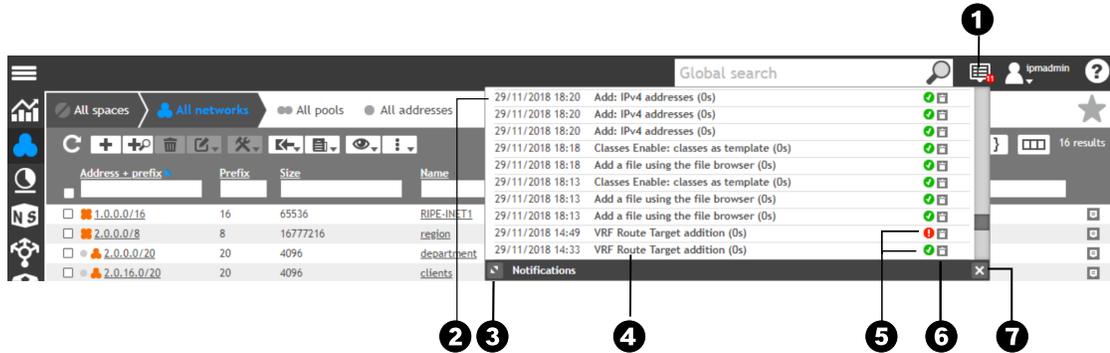


Figure 4.10. An example of results displayed in the window *Notifications*

- 1 In the top bar, click on **Notifications**, to open the window. The number in the red flag above the icon *Notifications* indicates the number of error messages in the window.
- 2 The notifications are listed in reverse chronological order.
- 3 This button allows to resize the notifications window.
- 4 The operation is displayed using the service it uses as follows: <action>:<resource it was performed on>.
- 5 The status is either if the operation was successful or if at least one error occurred. Any other icon indicates that the operation is ongoing.
- 6 This button allows to delete a notification. There is no warning before deletion.
- 7 This button allows you to close the window. Once you open it, the window remains open above the page you are currently on, until you close it.

To display and export notifications, click on the line of your choice .

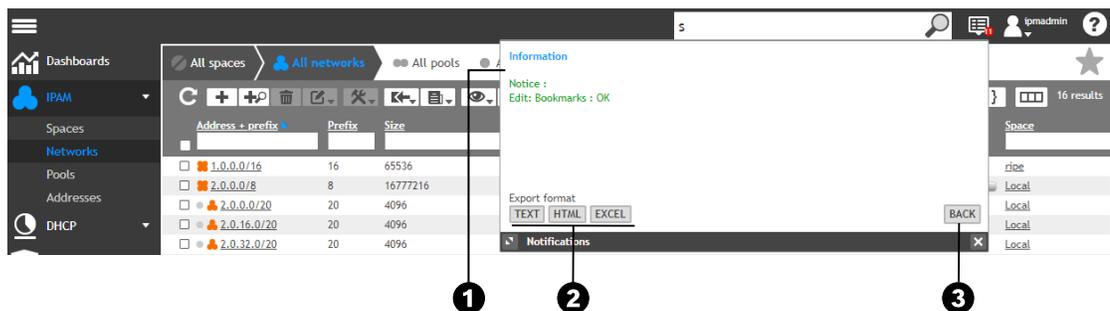


Figure 4.11. An example of a detailed result in the window *Notifications*

- 1 Once you clicked on a line, the type and full details of the report appear in the window.
- 2 Click on any of the available formats to export the final report.
- 3 This button allows to go back to the list of notifications.

To open the window **Notifications**

1. In the top bar, click on  **Notifications**. The window **Notifications** opens. The latest operations are displayed as follows: <date><hour:minute> <operation> <status> where *operation* contains the executed service as follows: <action>:<resource it was performed on>.
2. To display the report of an operation, click on the notification of your choice.
3. To go back to the list of **Notifications**, click on **BACK**.

To export the operation report

1. In the top bar, click on  **Notifications**. The window **Notifications** opens.
2. Click on the operation of your choice. The report opens.
3. Under the content of the report, click on **TEXT**, **HTML** or **EXCEL** to export the report in the corresponding format.

My Account

This menu is accessible from any page of SOLIDserver. It allows to:

- Access the connected user account configuration options via *My Settings*. For more details, refer to the section [Account Configuration](#).
- Change the connected user password.
- Access the page *My Quick Wizards*. For more details, refer to the section [Quick Wizards](#).
- Access the page *My Smart Folders*. For more details, refer to the chapter [Managing Smart Folders](#).
- Access the page *My Bookmarks*. For more details, refer to the section [Bookmarks](#).
- Access the page *My Gadgets*, that provides access to the pages *Gadgets Library* and *My Gadgets*. For more details, refer to the chapter [Managing Gadgets](#).
- Close your session via *Logout*.

Help

This menu allows to:

- Open SOLIDserver Administrator Guide, via *Administrator Guide*. The PDF file opens in a new tab of your browser. You can save it on your computer to access it even when you are no longer connected to the appliance.
- Open the Software License Agreement, via *License Agreement*.

Breadcrumb

The breadcrumb is a navigation bar. It is available on all the pages of SOLIDserver except in the module *Dashboards* and on most pages of the module *Administration*. It provides direct and hierarchical access to the module objects, allows you to filter data displayed on the listing pages and to access additional pages.

Browsing a Module Hierarchy

The breadcrumb gives access to the different objects of a module and allows to view several levels of hierarchy on a single line.



Figure 4.12. The breadcrumb

- ❶ The icons order respects the module internal hierarchy, from the highest level on the left down to the lowest level on the right. Here, the IPAM hierarchy: *All spaces*, *All networks*, *All pools* and *All addresses*.
- ❷ The icon and name displayed in blue show the page you are currently on, here the page *All networks*. The current level and its parent(s) are in dark gray.
- ❸ The lower levels are displayed in a light gray area. Here *All pools* or *All addresses*. You can access either directly.

Filtering Data

The breadcrumb allows to filter and browse the content of specific objects in each module.

Once the breadcrumb is filtered at a level:

- The page refreshes and longer displays the columns dedicated to the parent level.
- The breadcrumb no longer displays *All <objects>* but *<Object>: <object-name>*, and everything right of this filter only concerns the *object name*. For instance, if you filter the breadcrumb with a specific DNS server and zone, the breadcrumb provides access to both properties page.



Figure 4.13. Navigational information in the breadcrumb

- ❶ This display indicates that you are filtering data and browsing the content of the object mentioned. Here, you are displaying the views, zones and records of the DNS server *ns1.mycompany.com*.
- ❷ With the breadcrumb filtered at server level, clicking on *All views* opens the page but only displays the views of the server *ns1.mycompany.com*.
- ❸ With the breadcrumb filtered, the object level *All <objects>* is renamed *<Object>*. It still provides access to the page, here *Zone* opens the page *All zones* but only displays the zones of the server *ns1.mycompany.com* as the breadcrumb is filtered at server level. If it was not, it would list all existing zones regardless of their server.
- ❹ With the breadcrumb filtered, the object name is a link toward its properties page. Here, clicking on *mycompany.com* opens its properties page.
- ❺ With the breadcrumb filtered at zone level, clicking on *All RRs* opens the page but only displays the records of the zone *mycompany.com*.

Accessing Additional Pages

Some pages are not displayed by default in the breadcrumb. If some additional pages are available for an object, right to its name in the breadcrumb, you will see the chevron icon ».



Figure 4.14. Additional pages in the breadcrumb

- ❶ The chevron icon » indicates that additional pages are available. Click on the icon to display these pages in the breadcrumb.
- ❷ When the additional pages are displayed in the breadcrumb, you can click on « to hide them.
- ❸ All additional pages available are displayed to the right of the icon.

Menu

The menu is displayed on almost every listing page. Its content differs on each page and the options available vary from one user to another depending on the rights they are granted. For more details regarding rights, refer to the part [Rights Management](#).

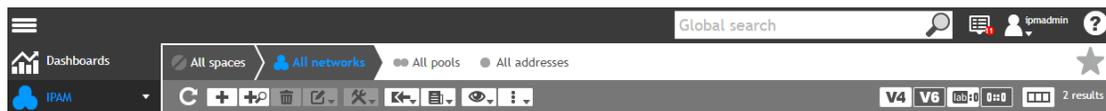


Figure 4.15. The menu

+ Add

This menu is available on all the **listing pages** when relevant.

In the IPAM, the DNS and Device Manager it allows to create the objects you manage on the page as well as the objects they can contain. For instance, in the module *DNS*, you can create DNS zones from the page *All servers*.

🗑 Delete

This menu is available on all the **listing pages** when relevant.

On some pages, the button 🗑 **Delete** is grayed out. You need to select the object that you want to delete in the list to be able to click on it.

✎ Edit

This menu is available on all the **listing pages** when relevant.

For example, it allows to synchronize/refresh servers and devices, manage access to the objects for groups of users or perform specific operations on the objects listed.

On the **properties page** of the objects, it allows to edit the panels on the page.

On some pages, the menu option  **Edit** is grayed out. You need to select the object that you want to edit in the list to have access to this menu.

Tools

This menu is available on all the **listing pages**. It allows to:

- Perform advanced operations on the objects listed.
- Perform *Expert* operations.

On the **properties pages**, this menu is available only for some specific advanced operations.

On some pages, the menu  **Tools** is grayed out. You need to select an object in the list to have access to this menu.

Import

This menu is available on all the **listing pages** when relevant.

In the IPAM, the DNS and Device Manager it allows to import the objects you manage on the page as well as the objects they can contain. For instance, you can import IP addresses from the page *All spaces*.

Report

This menu is available on the **listing pages** of all the modules when relevant. It allows to:

- *Export* the list in a CSV, HTML, XML, EXCEL and PDF file.
- Generate a report specific to the data listed in a HTML or PDF file.

This menu also allows to generate reports from some properties pages.

Alerts, Gadgets & Smart Folders

This menu is accessible from any listing page or properties page of SOLIDserver. It allows you to create alerts, gadgets and Smart folders.

For more details, refer to the chapters [Managing Gadgets](#), [Managing Alerts](#) and [Managing Smart Folders](#).

Extra Options

This menu is accessible from any listing page or properties page of SOLIDserver.

From **all pages** it allows to:

- Manage the listing templates created in the module, from the page *Listing Template Management*.
- Manage the import and export templates created in the module, from the page *Import/Export Templates Management*.

From the **listing pages**, it allows to:

- Display the configuration of the class *global* applied to the objects listed, via *Meta-data*. For more details, regarding classes refer to the chapter [Configuring Classes](#).
- To enable classes, via *Classes configuration*.
- Configure the advanced properties fields you want to enable or display in the addition/edition wizard of your IPAM, DNS or DHCP objects, via *Wizard customization*. For more details, refer to the chapter [Managing Advanced Properties](#).

Other Buttons

On the right-end side of the menu you can see other buttons.



Figure 4.16. Extra buttons on the right-end side of the menu

Listing Templates

This button allows to display, edit or create listing templates for the current page in a dedicated window. For more details, refer to the section [Customizing the List Layout](#).

V4 Switch to IPv4

This toggle buttons is displayed next to **V6** on the *IPAM* and *DHCP* pages that provide both IPv4 and IPv6 management. The gray button indicates the current version displayed. If you click on the white button, you display the other version.

V6 Switch to IPv6

This toggle buttons is displayed next to **V4** on the *IPAM* and *DHCP* pages that provide both IPv4 and IPv6 management. The gray button indicates the current version displayed. If you click on the white button, you display the other version.

Use IPv6 Labels

This button is displayed when the IPv6 toggle is active. It allows to use IPv6 labels. For more details, refer to the chapter [Managing IPv6 Labels](#).

Uncompress IPv6 addresses

This button is displayed when the IPv6 toggle is active. It allows to uncompress IPv6 addresses

Show/Hide

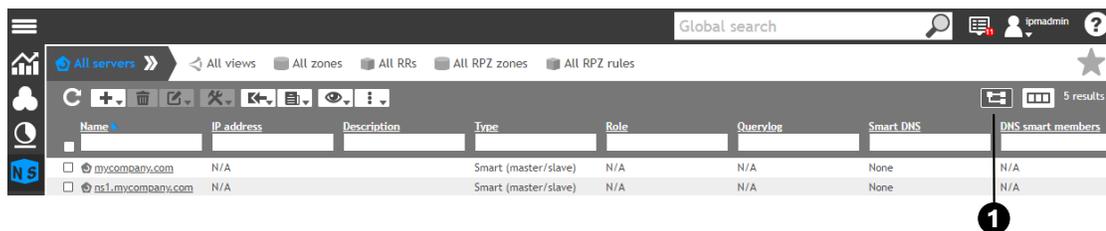


Figure 4.17. The button show/hide on the DNS page All servers

- 1 When the button is black, the physical servers are displayed (show).

In the DNS and DHCP, on the page *All servers*, this button allows to display/hide the physical servers managed via smart architectures under the *Name* of each smart architecture.

In the module Application, it allows to display/hide the deployed applications and traffic policies under the *Name* of each application.

In the module Guardian, it allows to display/hide the deployed policies under the *Name* of each policy.

Account Configuration

Each user has the possibility to edit their account once they are connected. From any page, they can use the menu **My Account** to edit their account preferences or they can go to the SOLID-server Main Dashboard to edit them from two default gadgets.

Configuring the User Display Settings

At any point, the connected user can configure their appliance preferences. These preferences include the display language, the number of lines displayed on listing pages or the time zone.

Any user can perform this configuration from the menu *My account*. The super user, *ipmadmin*, can also do it from the gadget *System Information*.

To configure the user settings

1. From the gadget **System Information**, in the section **Connected as**, click on the user name. The wizard **Configure user settings** opens.
2. In the field **List line count**, you can define how many entries (lines) you want to display on the listing pages. Keep in mind that the more you increase the number of lines displayed, the more SOLIDserver uses resources to display them.
3. In the drop-down list **List format**, you can set to alternate row colors every row, select *1-1* or every three rows, select *3-3*.
4. In the drop-down list **Display time in**, select *Local time* or *UTC-GMT*. All your services must be at the same time to prevent any management problems. Note that the *local time* is based on the time zone of your browser.
5. In the drop-down list **Date format**, select *mm/dd/yyyy* or *dd/mm/yyyy*.

6. In the drop-down list **Language**, you can set the interface language: *English, French, Spanish, German, Dutch, Chinese* or *Japanese*. By default, *English* is selected.
7. Click on to complete the operation. The report opens and closes. The *Main Dashboard* is visible again.

To configure the user settings from the gadget System Information

Only the super user *ipmadmin* can perform this operation.

1. From any page, in the top bar, select  **My account** > **My Settings**. The wizard **Configure User Settings** opens.
2. In the field **List line count**, you can define how many entries (lines) you want to display on the listing pages. Keep in mind that the more you increase the number of lines displayed, the more SOLIDserver uses resources to display them.
3. In the drop-down list **List format**, you can set to alternate row colors every row, select *1-1* or every three rows, select *3-3*.
4. In the drop-down list **Display time in**, select *Local time* or *UTC-GMT*. All your services must be at the same time to prevent any management problems. Note that the *local time* is based on the time zone of your browser.
5. In the drop-down list **Date format**, select *mm/dd/yyyy* or *dd/mm/yyyy*.
6. In the drop-down list **Language**, you can set the interface language: *English, French, Spanish, German, Dutch, Chinese* or *Japanese*. By default, *English* is selected.
7. Click on to complete the operation. The report opens and closes. You can see the display change on the Administration syslog page for instance.

Changing the Session Password

If you were granted sufficient rights, you can edit your local user password, the password used to connect to SOLIDserver.

Keep in mind that remote users cannot edit their password. Remote users come from a third party server or directory - AD, LDAP or RADIUS - and are authenticated via a dedicated rule. For more details, refer to the chapter [Managing Authentication Rules](#).

To change the password

Only local users can perform this action.

1. From any page, in the top bar, select  **My account** > **Change Password**. The wizard **Modify User Password** opens.
2. In the field **Previous password**, type in your old password.
3. In the fields **New password** and **Confirmation**, type in your new password.
4. Click on to complete the operation. The report opens and closes.

Listing Page

The listing pages are the most common pages within SOLIDserver.

- They allow to **sort** the data displayed on the page. When the name of the columns is underlined, it indicates that sorting is available and you can click on it.

- They allow to **filter** data. Use the search engine located under the column name to only display the objects matching your search.
- They allow to **manage** objects. The menu indicates what operations users can perform on the page. Depending on their rights, some options are not visible. For more details, refer to the part [Rights Management](#).
- They allow to **access the properties page** of the objects listed, when relevant. For more details, refer to the section [Properties Page](#).
- They provide a **contextual menu** with some options for the objects listed. For more details, refer to the section [Using the Contextual Menu](#).

The listing pages all share the same structure.

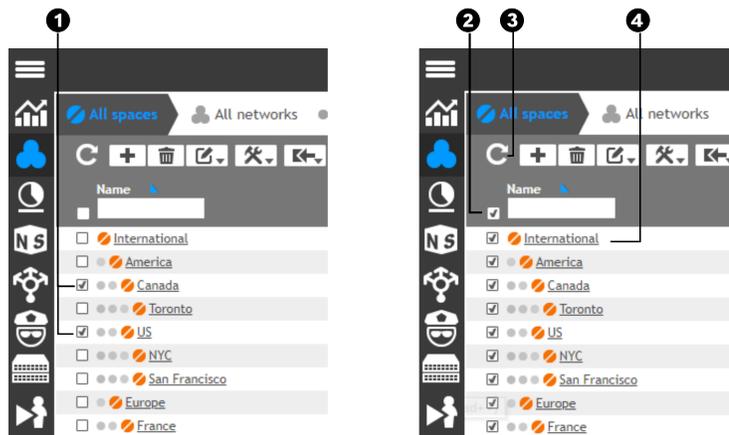


Figure 4.18. The buttons on all the listing pages

- 1 Each object **can be ticked and managed separately**. You can select a set of successive items using the key SHIFT on your keyboard.
- 2 **All the objects can be selected at once**. Above the list, you can tick the box left to the first column to select all the items listed, whether the list is filtered or not, it selects all the items counted in the *result*.
- 3 It allows you to refresh the listing page.
- 4 Within the data listed, all **the data underlined provides a link**. Depending on the page, it can be a filter to list the content of a container or open the properties page of an object. At the lowest level of a module hierarchy, it can perform specific operations (assign the object for instance) or open the edition wizard or properties page of the object.

There is **no limit to the number of objects** on a page but the more data you display, the more resources SOLIDserver uses to display them. To define the number of elements per page, refer to the procedure [To configure the user settings](#).

If there are more lines to be displayed than the number selected in your settings, the listing area **contains sub-pages** that allow to navigate within the object database. The GUI provides some key fields, buttons and areas to navigate within these pages. The following elements only appear when there are many elements on the same page:

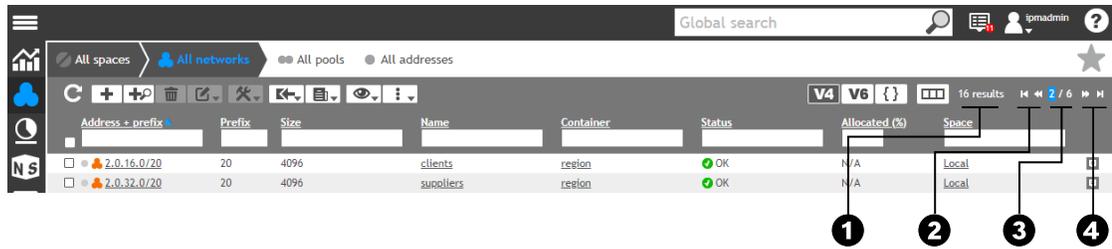


Figure 4.19. The buttons to navigate between sub-pages

- 1 The *results* indicates the total number of objects listed on all pages, if no filter is applied. If a filter is applied, it indicates the total number of objects matching your search.
- 2 These buttons allow you to display, respectively, the first and previous pages of data.
- 3 The number of pages on which data is listed. The number highlighted in blue displays which page you are currently on.
- 4 These buttons allow you to display, respectively, the next and last page of data.

Sorting the List

On all the listing pages, you can sort by ascending or descending order any column that has its name underlined.

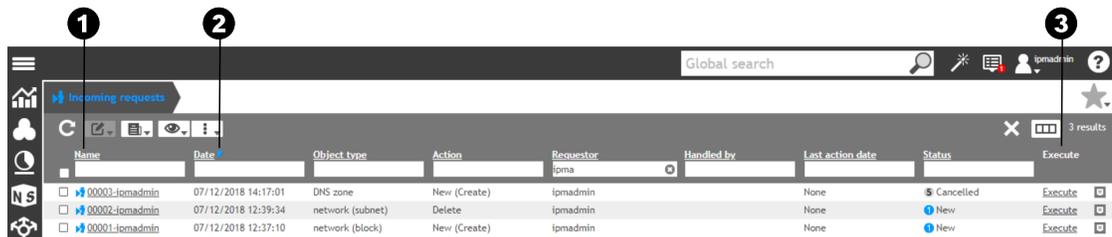


Figure 4.20. The columns' name can be used as a sorting tool

- 1 The underlined column names indicate that you can sort the data listed in direct/reverse alphabetical order.
- 2 This arrow indicates that the list is sorted through the column *Date* in ascending order.
- 3 The column names not underlined indicate that you cannot sort the column data.

To sort the list

1. Go to the module and listing page of your choice.
2. Click on any underlined column name. The list is sorted in descending order using the values in the column. If this is the column that sorts the list by default, it reverses the default order.
3. If you click on the name of the same column again, you change the sorting order.

Keep in mind that sorting a list:

- Follows the natural alphabetical order: first the digits are listed, then the letters, with no difference between lower and upper case.
- Can only be done one column at a time but can be used on a filtered list.

- If a column contains too much data it can no longer be sorted. In this case, you should find a  icon next to the column name.

Filtering the List

Almost all the data listed on the GUI pages can be filtered, using one or several columns.

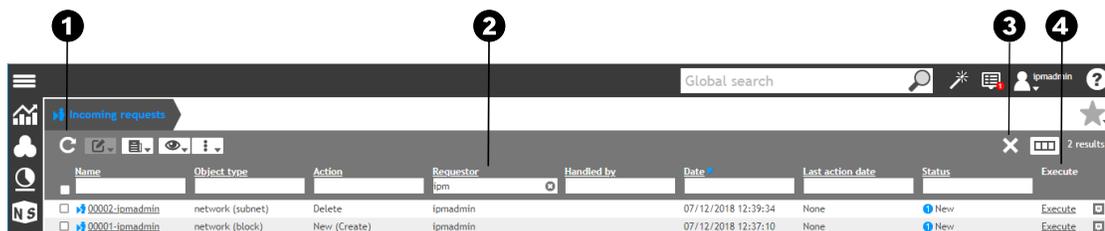


Figure 4.21. The column filtering tools

- 1 This button allows you to apply the filter using the data entered in the column search engine. You can also hit Enter to perform a filtered search.
- 2 The column search engine allows to type in the data you are looking for in the column. If this field is not visible, you cannot filter data via the column, you might only be able to sort.
- 3 This button allows to unset all the filters applied on the page, no matter how many were set.
- 4 Columns with no search field cannot be filtered.

You can filter the list using:

- The **operators** available in the search engine. They are all described in the table [Filtering operators](#).
- The **Filter constructor** of a column is available when you double-click in the search engine. It indicates the options accepted in each column and can also contain a list of possible statuses for the object or *Top occurrences* of the most represented values. In columns displaying times and dates, it contains specific operators that allow to filter dates or whole periods of time, based on UTC time. For more details, refer to the table [Date related filtering operators](#).

Using Filters

All the columns that can be filtered provide a search engine where you can type in or select, using the filter constructor, the values that you want to match or avoid.

There are filter constructors dedicated to the database recurring values, such as statuses, protocol versions, server types, etc.

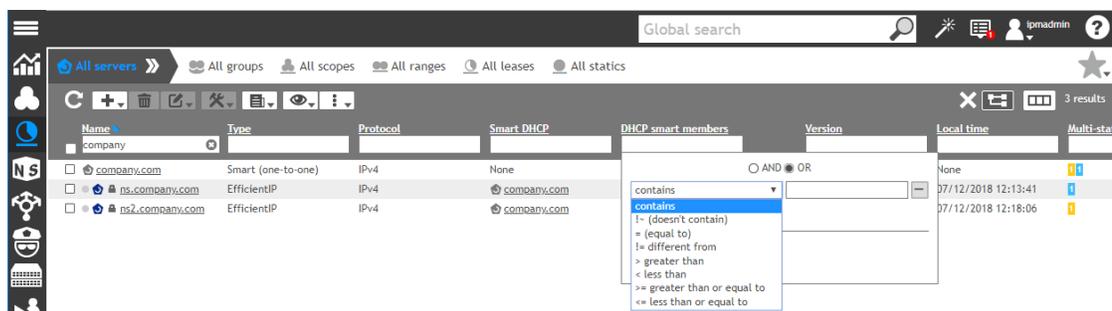


Figure 4.22. Filter constructor operators

To filter data from a specific column using operators

1. Go to the module and listing page of your choice.
2. In the search engine of any column, type in values or operators. For more details refer to the section [Accepted Operators for All Columns](#).
3. To perform the search, click on the button **C Refresh** or hit **Enter**. The list is filtered.
4. You can filter more columns if need be.

To filter data from a specific column using the filter constructor

1. Go to the module and listing page of your choice.
2. In the search engine of any column, double-click in the search engine. The filter constructor opens.
3. If you use several criteria, select the radio button **AND** or **OR**.
4. In the drop-down list, select a filter constructor. To see the list of available filter constructors, refer to the figure 1.17.
5. In the box on the right, enter the value you want to use to filter.
6. To use several criteria, click on **+**. A new line appears, follow the steps a and b to define the new criterion.

Repeat these actions for as many filters as needed.

7. To remove a line, click on **-** on the right-end side of the line.
8. In the search engine columns containing a specific data - like the *Status*, a server *Type*, the server *Sync* status, etc. - you can double-click in the search engine to display the list of existing values.
9. Click on **APPLY**. Only the data matching your search is displayed in the column.
10. You can filter more columns if need be.

For the recurring values in your database, filter constructors provide a box *Top occurrences*.

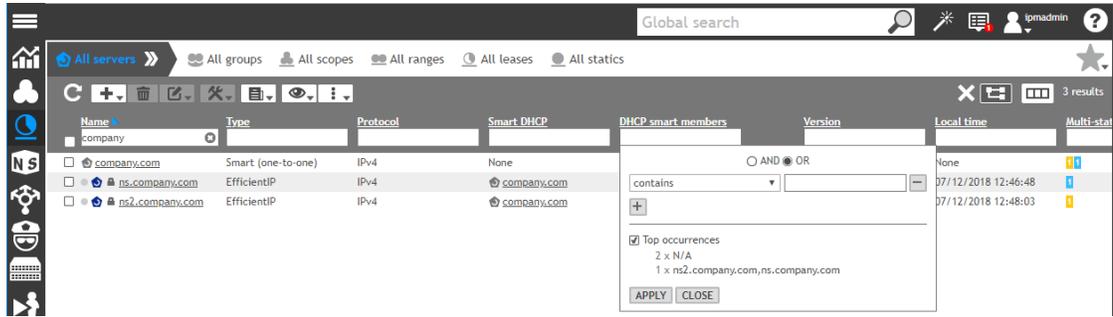


Figure 4.23. The top occurrences of DHCP physical servers on the page All servers

To filter data using the columns top occurrences

1. Go to the module and listing page of your choice.
2. In the search engine of any column, double-click in the search engine. The filter constructor opens.
3. If you use several criteria, select the radio button **AND** or **OR**.
4. In the drop-down list, select a filter constructor. To see the list of available filter constructors, refer to the figure 1.17.
5. Tick the box **Top occurrences**. The column top occurrences are displayed.
6. Click on a value (top occurrence). Its name is displayed in the field next to the filter you just selected.
7. To use several occurrences, click on **+**. A new line appears, follow the steps a and b to define the new criterion.

Repeat these actions for as many filters as needed.

8. To remove a line, click on **-** on the right-end side of the line.
9. Click on **APPLY**. Only the data matching your search is displayed in the column.
10. You can filter more columns if need be.

In the columns for time and date, there is a dedicated filter constructor.

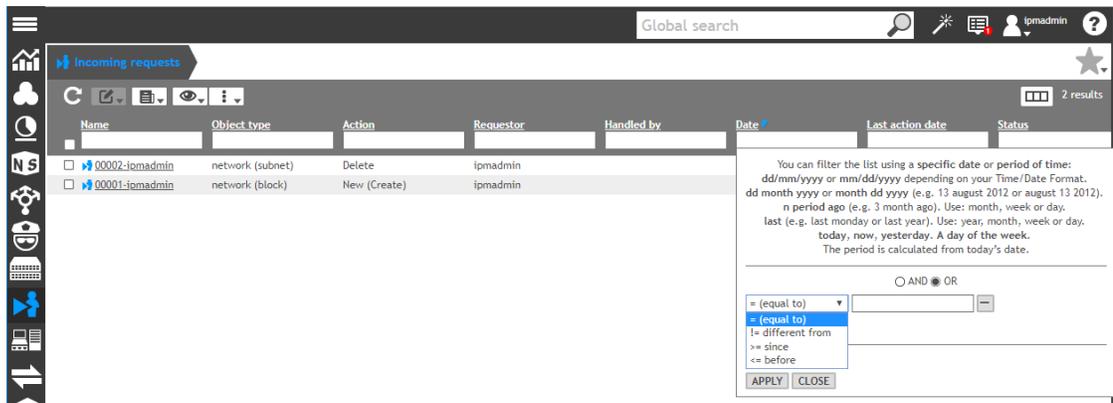


Figure 4.24. The dedicated time and date filters

To filter columns containing time and date using operators

1. Go to the module and listing page of your choice.
2. In the search engine of any column, type in values or operators. For more details refer to the section [Accepted Operators for Time & Date Columns](#).
3. Click on **APPLY**. Only the data matching your search is displayed in the column.
4. You can filter more columns if need be.

To filter columns containing time and date using the filter constructor

1. Go to the module and listing page of your choice.
2. In the search engine of any column, double-click in the search engine. The filter constructor opens. Depending on the resources listed on the page, it may provide a list of *Top occurrences*. This list provides an overview of the recurring values in the column (names, vendors...): the most used values.
3. If you use several criteria, select the radio button **AND** or **OR**.
4. In the drop-down list, select a filter constructor. To see the list of available filter constructors, refer to the figure 1.19.
5. In the box on the right, enter the value you want to use to filter.
6. To use several criteria, click on **+**. A new line appears, follow the steps a and b to define the new criterion.

Repeat these actions for as many filters as needed.

7. To remove a line, click on **-** on the right-end side of the line.
8. Click on **APPLY**. Only the data matching your search is displayed in the column.
9. You can filter more columns if need be.

Accepted Operators for All Columns

All the columns containing a search engine allow to filter data.

Note that these operators are not available in all search engine fields. For more details refer to the section [Using Filters](#)

For columns containing time and dates, refer to the section [Accepted Operators for Time & Date Columns](#).

You can also type in keywords in the search engine, the following operators are accepted:

Table 4.1. Filtering operators

Expression	Description
<i>string</i>	contains <i>string</i> .
<i>~string</i>	contains <i>string</i> .
<i>=string</i>	strictly equals <i>string</i> .
<i>>string</i>	greater than <i>string</i> .
<i><string</i>	less than <i>string</i> .
<i>>=string</i>	greater or equal to <i>string</i> .
<i><=string</i>	less than or equal to <i>string</i> .

Expression	Description
<code>!=string</code>	strictly different from <i>string</i> .
<code>!~string</code>	does not contain <i>string</i> .
<code>=#</code>	returns empty lines.
<code>!=# OR !~#</code>	returns lines containing data.
<code>*string</code>	ends with <i>string</i> .
<code>!~*string</code>	does not end with <i>string</i> .
<code>string*</code>	begins with <i>string</i> .
<code>!~string*</code>	does not begin with <i>string</i> .
<code>expression1 expression2</code>	<i>expression1</i> and <i>expression2</i> on the same line.
<code>expression1 & expression2</code>	<i>expression1</i> and <i>expression2</i> on the same line.
<code>expression1 expression2</code>	<i>expression1</i> or <i>expression2</i> in the same column: the line matching either data in the column is returned.

Keep in mind that the space between the operator and the string is accepted.

Accepted Operators for Time & Date Columns

On several pages, some columns are dedicated to time and/or date. If the column contains a search engine under the column name, it provides specific filtering operators.

The columns containing time and/or date also provide dedicated keywords that you can use to filter a list:

Table 4.2. Date related filtering operators

Expression	Description
<code>date</code>	Type in the date of your choice following the selected time and date format: <i>dd/mm/yyyy</i> or <i>mm/dd/yyyy</i> . You can also type in the month in full letters: <i>dd <month> yyyy</i> or <i><month> dd yyyy</i> .
<code>today</code>	The results only include data matching the date of the search.
<code>now</code>	The results only include data matching the time and date of the search.
<code>yesterday</code>	The results only include data matching the day before the date of the search.
<code>last</code>	The results only include all the data matching the <i>day</i> , <i>week</i> , <i>month</i> or <i>year</i> prior to date of the search.
<code>n period ago</code>	The results only include data matching the number <i>n</i> of <i>day</i> , <i>week</i> , <i>month</i> , <i>year</i> prior to date of the search.
<code>day</code>	Used with the keyword <i>ago</i> following the format <i><n period ago></i> , <i>day</i> or <i>days</i> allows to filter data based on a specific number of days prior to the date of the search.
<code>week</code>	Used with the keyword <i>ago</i> following the format <i><n period ago></i> , <i>week</i> or <i>weeks</i> allows to filter data based on a specific number of weeks prior to the date of the search.
<code>month</code>	Used with the keyword <i>ago</i> following the format <i><n period ago></i> , <i>month</i> or <i>months</i> allows to filter data based on a specific number of months prior to the date of the search.
<code>year</code>	Used with the keyword <i>ago</i> following the format <i><n period ago></i> , <i>year</i> or <i>years</i> allows to filter data based on a specific number of years prior to the date of the search.
<code>day of the week</code>	Any day of the week can provide a filter like <i>last <day-of-the-week></i> or <i>n <day-of-the-week> ago</i> . The column search engine is not case sensitive.

Removing Filters

Once you filtered columns, you can remove your filters following the procedures below.

To remove a filter on a column

1. Go to the module and listing page of your choice.
2. Once you filtered a column, its search engine contains the value you set and ✖.
3. Click on ✖ to remove the column filter.

To remove all filters

1. Go to the module and listing page of your choice.
2. Once you filtered a column, the cross ✖ appears on the right-end side of the menu.
3. Click on ✖ to remove the filters you set on the page.

If the page has columns filtered by default, they are still filtered.

Customizing the List Layout

Throughout SOLIDserver you can set customized column layouts, called listing templates, to display, hide and/or order the columns of your page on listing pages. Keep in mind that:

Listing templates are not available on all pages

- The pages that provide layout customization contain the menu  *Listing templates*. All the pages provide it, except *Analytics* (DHCP, DNS), *All configurations* (NetChange) and *All policies* (Guardian). In the module *Administration*, only the pages *Centralized Management*, *Users* and *Groups* provide it.
- The pages that do not provide it display all the existing columns.

On all pages providing listing templates, a default one is available

- By default, any page providing layout customization displays the listing template *default*.
- You can edit it and if you do so, the modifications are immediately displayed, but we strongly recommend creating new templates instead because editing the *default* template overwrites the initial page settings and prevents you from getting back to the original page layout.
- It is impossible to rename or delete the templates *default*. You can only edit the list of columns it contains.

Listing templates have to be displayed

- Creating a listing template does not display it on the page. Once created, it is available in the menu  *Listing templates* where you can select it.

The management of listing templates depends on the user permissions

- Any user can display an existing listing template on the pages that provide it.
- Only a user of the group *admin* can add, edit and remove listing templates. Note that editing a listing template affects the layout for all the users that display it.
- Only a user of the group *admin* can move the columns on a page. This operation updates the listing template displayed.

All modules provide a management page for listing templates

- In each module, the page *Listing Templates Management* allows to manage the listing templates for all the module pages that provide layout customization. This page is accessible from any page of the module through the menu **⋮** *Extra options* > *Listing Templates Management*.
- On this page, the *default* listing templates are only available after you have added other listing templates.

Classes can help customizing the layout

- You can include columns displaying class parameters in the listing templates, these columns are all named *Class param: <object label>*. For more details regarding classes, refer to the section [Configuring Classes](#).
- You can automatically display listing templates based on the class applied to the container of the object listed via the **automatic templates**. Note that the automatic templates are not available at the highest level of a module hierarchy. For more details on how to apply a class, refer to the section [Configuring Classes](#).

Displaying Listing Templates

All users can display one of the existing listing templates.

To display a listing template

1. Go to a listing page that has at least one custom listing template.
2. On the right-end side of the menu, click on **☰ Listing Templates**. The window opens. The list of columns is displayed. The columns currently displayed are ticked in this list.
3. In the drop-down list *Displayed listing template*, select *<your-template>*. The page refreshes. Only the columns of the template are displayed.

To display an automatic listing template

1. Go to a listing page that has at least one automatic template.
2. On the right-end side of the menu, click on **☰ Listing Templates**. The window opens. The list of columns of the current template is displayed.
3. In the drop-down list *Displayed listing template*, select *Automatic template*. Only the columns of the template are displayed.
4. Click on the name of a container configured with one of the classes selected in the automatic template to display the objects it contains. The page refreshes and only the columns configured in the template are displayed.

Note that if your administrator has configured several templates with the same class, the *Automatic template* option displays the most recent one.

Adding Listing Templates

Only users of the group *admin* can perform this operation.

You can add listing templates and automatic templates from listing pages or the page *Listing templates management*.

Note that when you add a new listing template, it is not automatically displayed.

To add a listing template from a listing page

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, click on **Listing templates**. The window opens.
3. Click on **MORE OPTIONS**. The wizard opens.
4. In the drop-down list **Action**, select *New template*. A field appears.
5. In the field **Name**, type in the template name.
6. Click on **NEXT**. The page **<Objects> lists configuration** opens.
7. Via the lists **Hidden columns** and **Displayed columns**, configure your listing template. Each list contains the same columns as the template *default*.
 - To add a column to the template, select it in the list **Hidden columns** and click on . The column is moved to the *Displayed columns*.
 - To remove a column from the template, select it in the list **Displayed columns** and click on . The column is moved to the *Hidden columns*.
 - To set the order of the columns, select them one by one in the list **Displayed columns** and click on or .

If you are adding a template at the highest level of a module, go to the step 10.

8. Click on **NEXT**. The last page of the wizard opens.
9. In the drop-down list **Parent level**, *None* is selected.
10. Click on **OK** to complete the operation. The report opens and closes.
11. Display the template. For more details, refer to the section [Displaying Listing Templates](#).

To add a listing template from the page Listing templates management

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, select **Extra options > Listing templates management**. The page opens.
3. In the panel of your choice, click on **ADD**. The wizard **Template Selection** opens.
4. In the drop-down list **Action**, select *New template*. A field appears.
5. In the field **Name**, type in the template name.
6. Click on **NEXT**. The page **<Objects> lists configuration** opens.
7. Via the lists **Hidden columns** and **Displayed columns**, configure your listing template. Each list contains the same columns than the template *default*.
 - To add a column to the template, select it in the list **Hidden columns** and click on . The column is moved to the *Displayed columns*.
 - To remove a column from the template, select it in the list **Displayed columns** and click on . The column is moved to the *Hidden columns*.
 - To set the order of the columns, select them one by one in the list **Displayed columns** and click on or .

If you are adding a template at the highest level of a module, go to the step 10.

8. Click on **NEXT**. The last page of the wizard opens.
9. In the drop-down list **Parent level**, *None* is selected.
10. Click on **OK** to complete the operation. The report opens and closes.
11. Display the template. For more details, refer to the section [Displaying Listing Templates](#).

Adding Automatic Templates

Only users of the group *admin* can perform this operation.

An automatic template is a listing template that is automatically displayed when one or several classes, defined in the template, of the current object container are selected.

You can add automatic templates from listing pages or the page *Listing templates management*.

To add an automatic template

Only users of the group *admin* can perform this operation.

You can add a template for page listing objects that have a container.

1. Go to the listing page of your choice.
2. In the menu, click on **Listing templates**. The window opens.
3. Click on **MORE OPTIONS**. The wizard **Template Selection** opens.
4. In the drop-down list **Action**, select *New template*. A new field appears.
5. In the field **Name**, type in the template name.
6. Click on **NEXT**. The page **<Objects> lists configuration** opens.
7. Via the lists **Hidden columns** and **Displayed columns**, configure your listing template. Each list contains the same columns than the template *default*.
 - To add a column to the template, select it in the list **Hidden columns** and click on **+**. The column is moved to the *Displayed columns*.
 - To remove a column from the template, select it in the list **Displayed columns** and click on **-**. The column is moved to the *Hidden columns*.
 - To set the order of the columns, select them one by one in the list **Displayed columns** and click on **▲** or **▼**.
8. Click on **NEXT**. The last page of the wizard opens.
9. In the drop-down list **Parent level**, select the level for which the class should be applied. Two lists appear.
10. Select the classes that automatically display the template:
 - To add a class, select it in the list **Available classes** and click on **+**. It is moved to the list *Selected classes*.
 - To remove a class, select it in the list **Selected classes** and click on **-**. It is moved to the list *Available classes*.

11. Click on **OK** to complete the operation. The report opens and closes.
12. Display the template. For more details, refer to the section [Displaying Listing Templates](#).

Editing Listing Templates

Only users of the group *admin* can perform this operation.

You can edit any type of listing template from listing pages or the page *Listing templates management*.

Note that moving the columns of a page automatically edits the current listing template, but you cannot remove a column from the template that way.

To move a column

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. Hover over the header of the column you want to move, next to its name.
3. Click and drag the column.
4. Drop it where you want. The new location of the column updates the current listing template. If you click on **Listing template**, in the window, you can see that the order of the columns is updated.

To edit a listing template

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, click on **Listing templates**. The window opens.
3. Click on **MORE OPTIONS**. The wizard **Template Selection** opens.
4. In the drop-down list **Action**, select *Edit: <your-template>*.
5. Click on **NEXT**. The page **<Objects> lists configuration** opens.
6. In the lists **Hidden columns** and **Displayed columns**, edit the columns according to your needs.

If you are editing a template at the highest level of the module, go to the step 9.

7. Click on **NEXT**. The last page of the wizard opens.
8. In the drop-down list **Parent level**, select a different value if need be. For more details, refer to the procedure [To add an automatic template](#).
9. Click on **OK** to complete the operation. The report opens and closes.
10. Display the template. For more details, refer to the section [Displaying Listing Templates](#).

To edit a listing template from the page Listing templates management

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, select **⋮ > Extra options > Listing templates management**. The page opens.

3. In the panel of your choice, select the template that you want to edit.
4. Click on **EDIT**. The wizard opens.
5. In the lists **Hidden columns** and **Displayed columns**, edit the columns according to your needs.

If you are editing the template at the highest level of the module, go to the step 8.

6. Click on **NEXT**. The last page of the wizard opens.
7. In the drop-down list **Parent level**, select a different value if need be. For more details, refer to the procedure [To add an automatic template](#).
8. Click on **OK** to complete the operation. The report opens and closes.
9. Display the template. For more details, refer to the section [Displaying Listing Templates](#).

Renaming Listing Templates

Only users of the group *admin* can perform this operation.

You can only rename listing templates, only from the page *Listing templates management*.

Note that if you rename a listing template while it is displayed on the page, you need to select it again in the listing page.

To rename a listing template

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, select **⋮ > Extra options > Listing templates management**. The page opens.
3. In the panel of your choice, select the template you want to rename.
4. Click on **RENAME**. The wizard opens.
5. In the field **New Name**, type in the template new name.
6. Click on **OK** to complete the operation. The report opens and closes. The page is visible again, the new name is displayed.

Deleting Listing Templates

You can only delete listing templates from the page *Listing templates management*.

Note that you can delete a listing template even if it is being displayed.

To delete a listing template

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, select **⋮ > Extra options > Listing templates management**. The page opens.
3. In the panel of your choice, select the template you want to delete.
4. Click on **DELETE**. The wizard opens.
5. Click on **OK** to complete the operation. The report opens and closes. The page is visible again, the template is no longer listed.

If you delete the listing template that is currently displayed, the automatic template is displayed.

Displaying the Contextual Menu

On every listing page, SOLIDserver provides a right-click contextual menu for each object listed. It provides a few shortcuts and some minimal options also available on the properties page of the object. The content of the menu differs on every page and column.

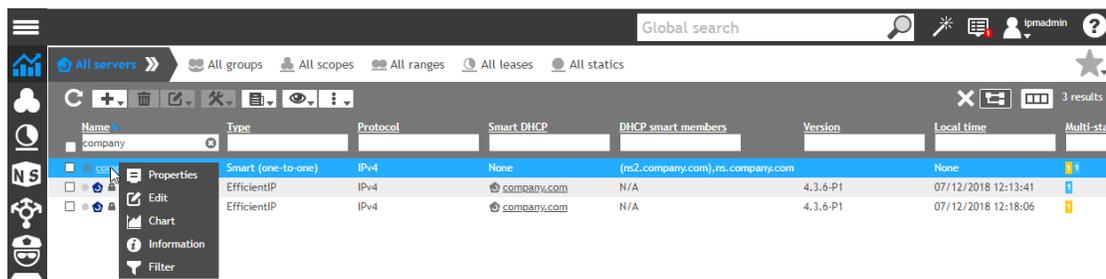


Figure 4.25. The full contextual menu

At most the menu offers 5 buttons:

- **Properties** is a shortcut to the properties page of the object.
- **Edit** is a shortcut to open the object's edition wizard. It opens the same wizard as the button *EDIT* of the panel Main properties on the properties page.
- **Chart** displays a statistics chart: this chart is also available on the properties page of the object.
- **Information** displays a table containing the basic information of the object. This information is also displayed in the panel Main properties on its properties page.
- **Filter** allows to filter the list using the value you right-clicked on.

For more details regarding the options available on the properties page, refer to the section [Properties Page](#).

Understanding the Column Multi-status

The column **Multi-status** returns messages if the current configuration of an object is worth mentioning or very specific. When the column returns information, a colored square containing a number appears on the line of the object.

The messages returned by this column do not always reflect configuration errors for the object. For instance, in the DNS, the Multi-status message *61006: Server type incompatible with Hybrid* indicates that the server in question cannot be switched to Hybrid DNS, it is probably managing authoritative and recursive zones; it does not mean that the server is not running properly or is misconfigured.

The column is displayed by default on some pages of the modules DNS and DHCP. You can display it on the other pages. For more details, refer to the section [Customizing the List Layout](#).

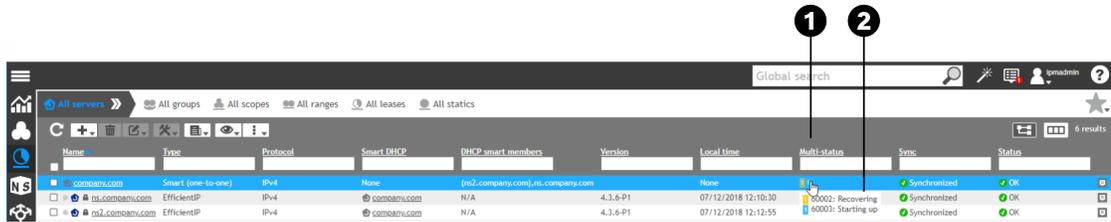


Figure 4.26. Example of Multi-status messages on the DHCP page All servers

- 1 The column **Multi-status** returns colored squares containing messages. The color indicates the **severity** of the message. In each square, a number indicates how many messages match the severity. In this example, the colored squares indicates that there is 1 error (yellow) message and 1 warning (blue) message returned for the smart architecture and its content.
- 2 Hovering over a square opens the window containing the message(s) of the object. In this example, one physical server returns the message 60002 and the other one the message 60003.

The column provides messages divided into 6 levels of severity. Each one provides useful status and state information regarding the object or the configuration within the module.

Table 4.3. Multi-Status severity levels

Severity	Color	Description
Emergency	Red	The object configuration prevents the system from running properly. Action is required.
Critical	Orange	The object configuration is in critical conditions. Immediate action is recommended.
Error	Yellow	The object configuration failed at some level. Action is recommended.
Warning	Blue	The object configuration triggers error messages if no action is taken. Action to be taken at your discretion.
Notice	Light blue	The object configuration is normal but undergoing events that might trigger errors. No immediate action is required.
Informational	Gray	The object configuration is normal, operational messages (might inform you about potential incompatibilities with other modules, etc). No action is required.

Each message and level of severity is specific to each object. The number of the messages are distributes among the modules: all DHCP messages start with 60000, all DNS messages with 61000... For more details, refer to the appendix [Multi-Status Messages](#).

Displaying IPv6 Labels

IPv6 Labels allow to customize the display of IPv6 addresses managed on a set of pages of the modules IPAM, DHCP, NetChange and Application.

They are displayed above start addresses to highlight IPv6 containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Address + prefix	Prefix	Name	Space	Container	Status
EUR::/16	16	Europe	Local	N/A	OK
CAN::/16	16	Canada	Local	N/A	OK
US::/15	15	US	Local	N/A	OK
US:0:0:1::/64	64	East Coast	Local	US	OK
US:0:0:1::/64	64	NYC	Local	East Coast	OK
US:0:0:2::/64	64	Midwest	Local	US	OK
US:0:0:3::/64	64	West Coast	Local	US	OK

Figure 4.27. Example of IPv6 labels used to highlight a geographical distribution in the IPAM

Properties Page

The properties page gathers all the information regarding an object.

- It is accessible from the listing pages. If the listing page already gathers all the information regarding the object managed, the object has no properties page.
- It allows to configure or edit object configurations. Some options are only available on the properties page of the object.
- It distributes all the information among panels. The panel Main properties contains the most general information. All the other panels contain more specific data. The panels that can be edited contain the button EDIT.
- It contains the panel *Audit* in the modules IPAM, DNS and DHCP that indicates every change made on the object, by whom and when.
- It contains common panels within a module and levels in the hierarchy.

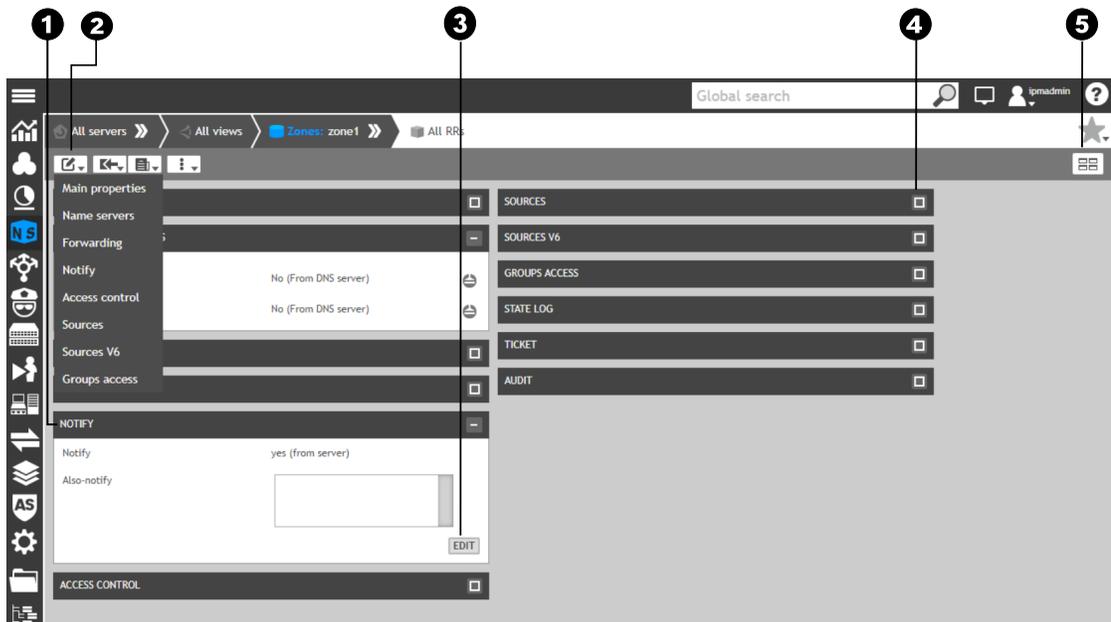


Figure 4.28. An example of properties page

- 1 This is what an open panel looks like. In this example, the screenshot displays the default display of a zone properties page.

- 2 In the menu, the menu  **Edit** presents the editable panels in one list.
- 3 The button  is present at the bottom of each and every open panel that you can edit. It opens an edition wizard of the object parameter.
- 4 This is what a closed panel looks like: only the panel name is visible. You can open it via the button .
- 5 The button  **Expand/Collapse All** allows to open or close at once all the panels of the page.

Some panels are specific to the object, others are available across the modules:

- **Main properties:** provides an overview of the main information regarding the object.
- **Advanced properties:** displays the advanced properties configuration of the object and the level the property was inherited from.

This panel is available in the modules IPAM, DHCP and DNS for the resources that can be configured with advanced properties, for more details refer to the chapter [Managing Advanced Properties](#).

- **Audit:** logs all the changes carried out on the object by the connected user over time.

This panel is available on all properties pages of the modules IPAM, DHCP and DNS; except for DHCP servers, groups, scopes, leases and DNS views and RRs.

If the user belongs to a group with access to the modifications of all users, it displays all the operations ever performed on the object. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).

- **Groups access:** displays all the groups that have the object as resource. Under each group name are listed the rights granted over the resource and its content.

This panel is available in the modules IPAM, DHCP and DNS on all properties page, except at the lowest level of each hierarchy and for DHCP groups and ranges.

In the module Administration, it is available on the properties page of users, except *ipmadmin*. Only the name of the group is displayed but not the rights granted to the group. For more details, refer to the part [Rights Management](#).

Charts

In SOLIDserver, there are two types of charts:

1. **Time-based charts:** they take the form of line or stacked-line charts. They offer a set of options that allow you to highlight, zoom in or display a specific period. All these functionalities are detailed below.
2. **Instant charts:** they take the form of pie and bar charts. They do not offer the same options: on pie charts you cannot select a period; on bar charts you cannot highlight data.

Charts are available on the properties pages of some objects, on the page *System statistics*, in gadgets or when you generate some reports. For more details, refer to the sections [Properties Page](#), [Monitoring the Appliance Statistics](#), [Managing Gadgets](#) and [Managing Reports](#).

In the image below the options are detailed on a time-based chart but can also be available on instant charts.

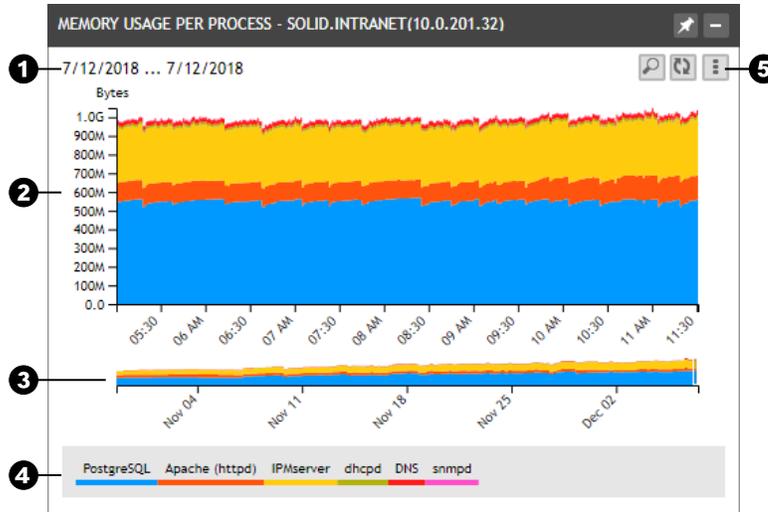


Figure 4.29. Overview of a time-based chart

- ❶ The start and end dates of the data displayed, it matches the period selected in the timeline and affects the scale of the x-axis. By default, each chart displays the *last 3 hours*.
- ❷ The data retrieved is represented in a line or stacked chart. The y-axis indicates the unit, axis scale and unit prefix depend on the period selected and maximum value displayed. Following the standard ISO 80000-1, all the y-axis units can have no prefix or any SI prefix such as: *m* (milli), *k* (kilo) or *M* (mega).
- ❸ The timeline of any time-based chart: the overall period of data available. The period displayed is highlighted in gray. By default, it displays a maximum of 365 days. To edit it, refer to the section [Editing the Number of Days Available on the Timeline](#).
- ❹ The legend of the chart. Each set of data has a name and a dedicated color.
- ❺ The display options allow to open the chart in a pop-up window with the button , to refresh the data with  or select a period with  that opens a drop-down list that allows to select between *Current hour*, *Last 3 hours*, *Day*, *Week*, *Month* or *Year*. It also provides access to  to select a specific date with.

Setting the Period to Display

The period selected in the timeline can be extended, reduced or moved.

If you hover over the edge of the gray area, the mouse pointer turns into a two-sided arrow that you can move left or right according to your needs.



Figure 4.30. Extending or reducing the period selected on the timeline

This period, represented by a gray area can also be dragged to select a different period.

If you hover over the gray area, the mouse pointer turns into a four-sided arrow.

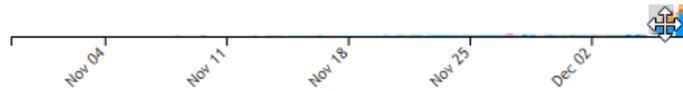


Figure 4.31. Sliding along the timeline

In addition, you can click and drag on the timeline to select a period directly.

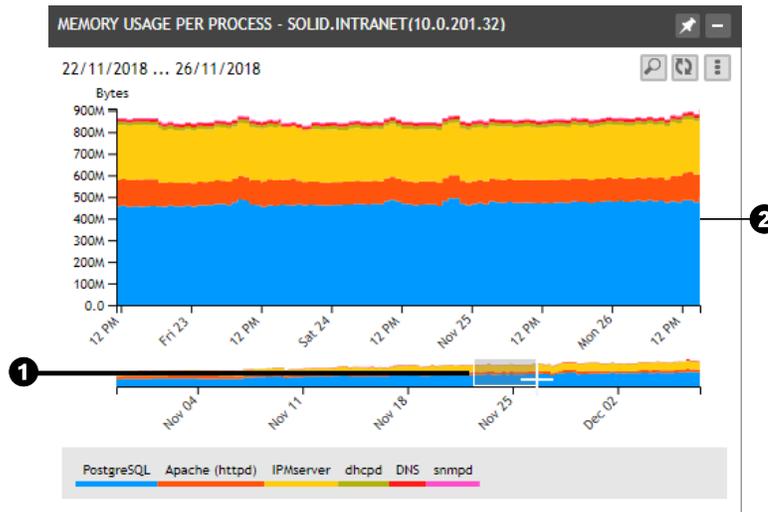


Figure 4.32. Zooming in on the timeline of a chart

- ❶ Within the timeline, with a left-click of the mouse over a white area, you can select the period that suits your needs. The pointer changes from an arrow to a cross.
- ❷ Once you release the mouse, the data displayed in the chart, x-axis points of reference and y-axis scale adjust accordingly.

Displaying Some Data and Zooming in

You can zoom in and out directly on the chart. The period selected when you zoom automatically updates the x-axis of the chart and the period selected in the timeline.

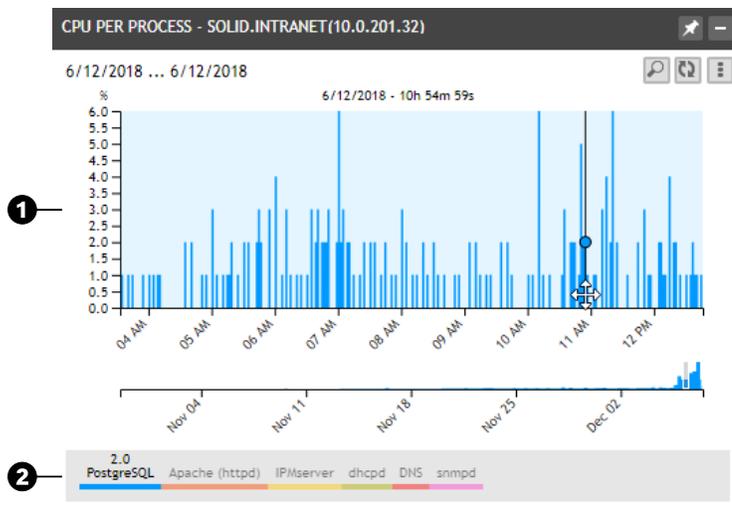


Figure 4.33. Zooming in on specific data

- 1 When you hover over the chart, the arrow pointer turns into a four-sided arrow and the background turns blue to indicate that you can interact with the data. You can display the data date and time above the chart and the value at that time is displayed in the legend of the chart.

A black vertical line indicates where you are on the chart, each measurements focus is symbolized by a circle in the color of the element displayed. Using the scroll wheel on the mouse, you can zoom in and out.

- 2 You can click on the elements of the chart legend to display or hide them from the chart. When you are browsing the chart, the value of each measurement on the line is indicated above each element of the legend.

Editing the Number of Days Available in the Timeline

The timeline of every time-based chart indicates for how long they display data. At first you only have hours and then days, weeks, months and finally a year. The default period of data displayed in the timeline is one year, 365 days.

A registry database entry allows to change that default maximum value. Once you edit it, all time-based charts refresh and display only the data retrieved over the period, number of days, that you set.

To edit the registry key that sets the number of days displayed in the timeline

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in *module.graph.default.period*.
4. Hit **Enter**. Only this key is listed.
5. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
6. In the field **Value**, type in the value of your choice, in days. The default value is 365.

- Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Wizards

Within SOLIDserver every operation - an object addition, edition, configuration, deletion - is performed via a wizard. All the modules share a common wizard structure, the fields and/or buttons that it contains depend on each operation. The title of the wizard specifies the ongoing operation.

In addition to the wizards, SOLIDserver uses pop-up windows: when there are configuration errors or when you select too many or not enough objects from a list before performing operations via the menu. However, some pages, like the Administration pages *Groups* and *Class Studio* or the IPAM page *All addresses*, use pop-up windows. Therefore, **to use SOLIDserver to the best of its potential, make sure your Internet browser is not configured to block pop-up windows.**

All the wizards share a common structure detailed in the sections below.

Common Structure of the Wizards

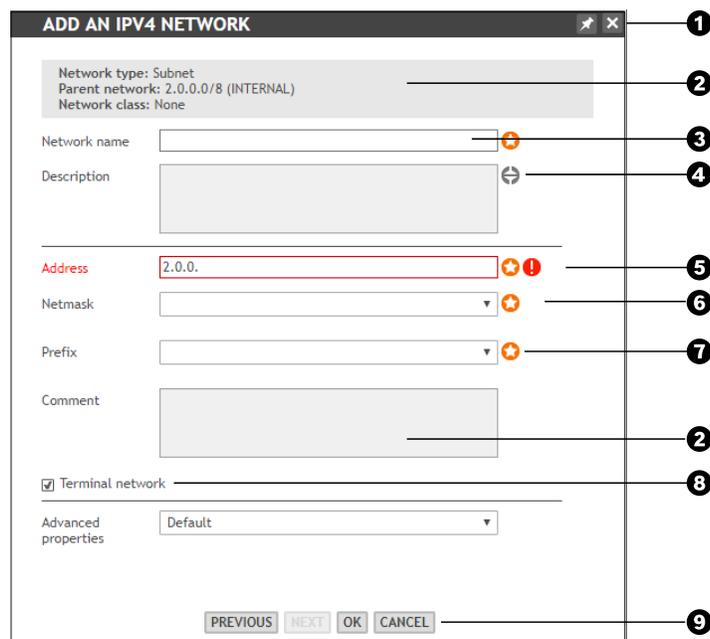


Figure 4.34. The common structure of the wizards

- The drag bar contains the title, a pushpin to save the wizard and a cross to close the wizard without saving any changes. For more details regarding how to save a wizard, refer to the section [Quick Wizards](#).
- The gray areas are informational sections. The first one is a location reminder: it indicates the container you are in and its basic information. The other sections are read-only sections, like the field *Comment* that sums up the current configuration, or informational messages to guide you during the configuration.
- The input fields are the most commonly used. Their border changes color in case of misconfiguration.

- 4 The button Set/Propagate allows you to configure the inheritance or propagation properties of the value in the field. If the value is inherited, the field background is gray and cannot be edited. For more details, refer to the section [Setting Advanced Properties](#).
- 5 The fields name and border turn red if there is a syntax error. In addition, the exclamation mark icon is displayed and you cannot save the configuration.
- 6 The star icon indicates that a field/parameter/option is mandatory. If you leave the field blank, you cannot save your configuration. If the field or drop-down list has a default value that you do not change, it is selected and applied when you commit the configuration.
- 7 The field drop-down list contains a down arrow to indicate you have several values to choose from.
- 8 The box is present on many wizards to set specific parameters. Ticking it usually reloads the wizard and allows to set specific parameters in extra fields.
- 9 The navigation buttons of the wizard. The button **OK** indicates that you are on the last page of the wizard. Clicking on it saves and applies your configuration. The button **CLOSE** closes the wizard without saving the configuration or changes applied.

For more specific configurations, the wizards embed extra information icons. These icons open a window containing more detailed information to help with a thorough configuration of the object.

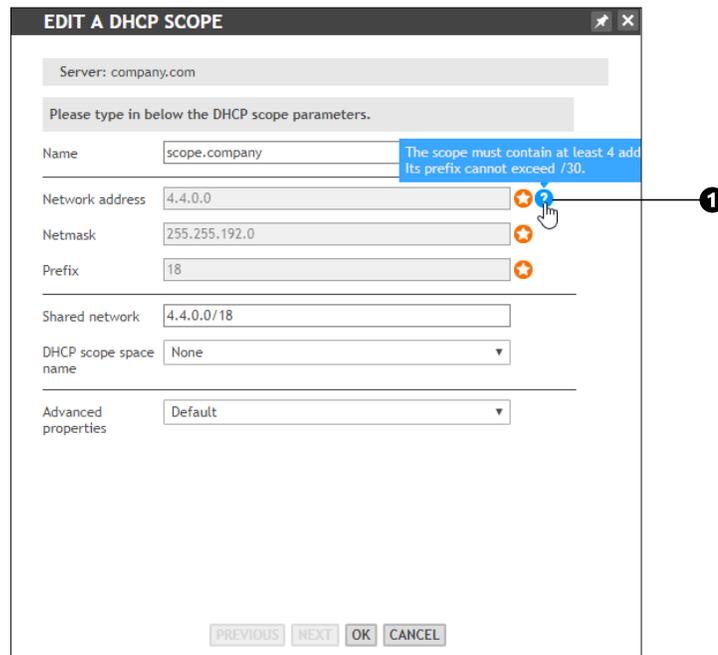


Figure 4.35. An example of help information buttons in the wizards

- 1 The question mark icon indicates extra details regarding a field. Hover over it to open the field configuration help.

Messages in the Wizards

A number of wizards include warning and information messages that you should take into account before saving your changes or configuration.

For instance, all the object deletion wizards contain a warning message to make you confirm the deletion or to provide extra information regarding the consequences of the deletion.

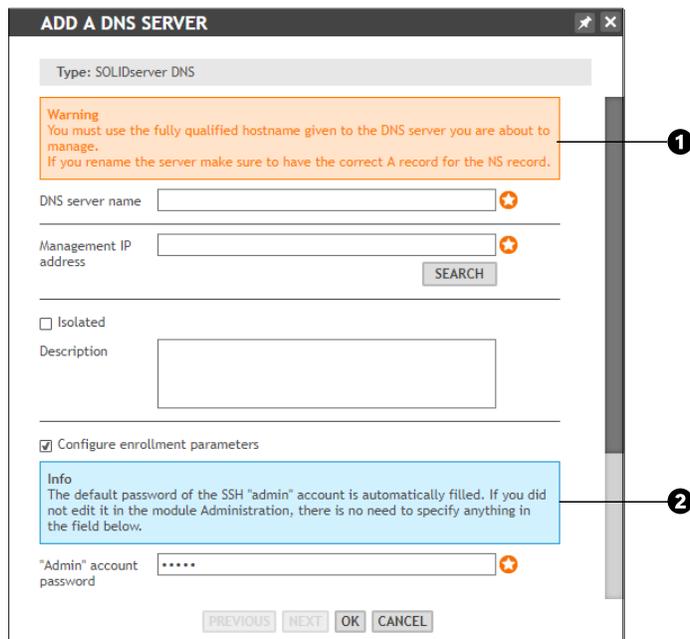


Figure 4.36. An example of messages in the wizards

- 1 Warning messages are displayed in orange. Some specific required values that cannot be directly verified by the wizard are introduced by warning messages.
- 2 Information messages are displayed in blue.

Configuration Lists in the Wizards

A number of wizards provide configurations lists to manage data. They contain two lists that gather all available data and allow you to choose values from a list or set up your own list. They usually go in pairs: Available/Selected or Hidden/Displayed.

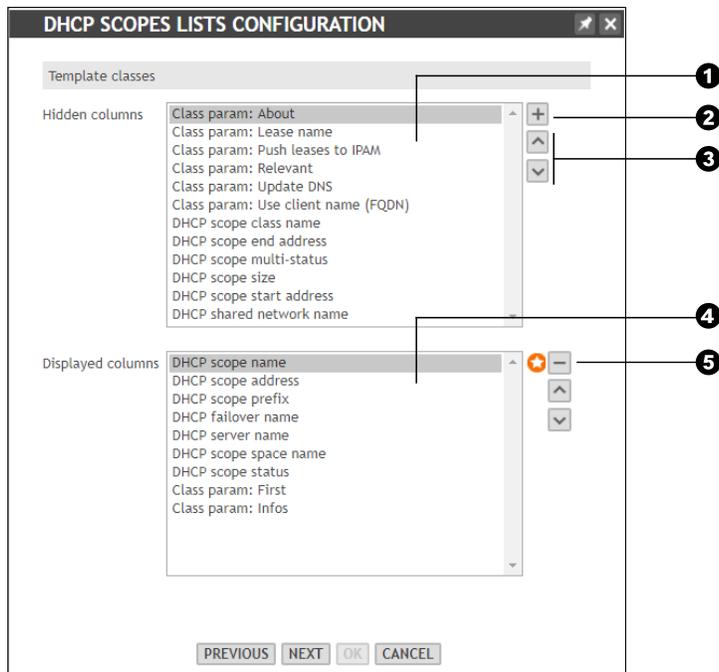


Figure 4.37. An example of configuration lists in a wizard

- ❶ This list displays all the available columns that are not yet configured in the listing template. In this example, the columns that could be displayed on the page *All scopes*.
- ❷ Once you have selected a value in the list *Hidden column*, click on **+** to move it to the list *Displayed columns* and include them into the listing template.
- ❸ These buttons allow to order the list entries. Select them one at a time and move them up **▲** or down **▼** until the order suits your needs.
- ❹ This list displays all the columns that are part of the listing template.
- ❺ You can remove any column from the listing template. Select them one at a time in the list *Displayed columns* and click on **-** to move them to the list *Hidden columns*.

The module Administration provides a set of wizards where you can set up and edit multiple entries in a single list.

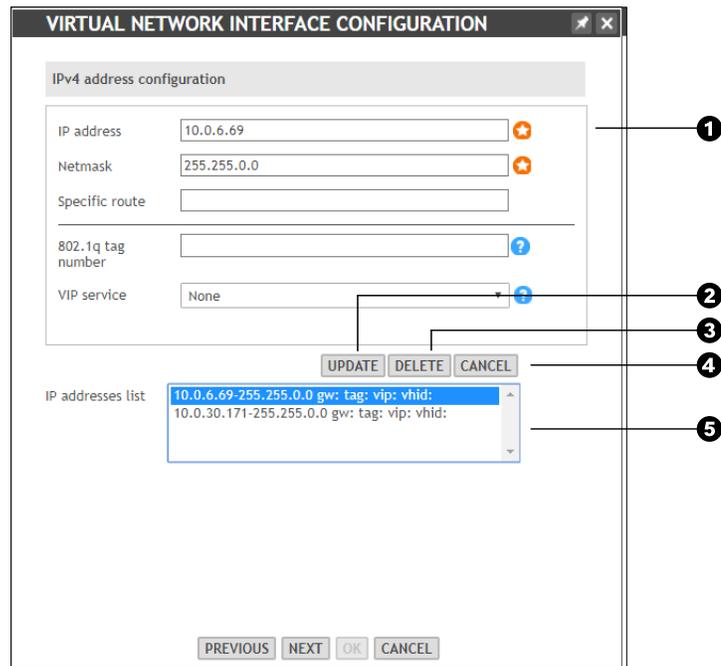


Figure 4.38. An example of data edition in a wizard providing entries management in one list

- ❶ Once you have selected an entry in the list at the bottom of the wizard page, in this example *IP addresses list*, its configuration details appear in these fields. You can edit any white field.
- ❷ Click on **UPDATE** to save your modifications and overwrite the former configuration and follow the wizard to commit your changes.
- ❸ Click on **DELETE** to delete the selected configuration entry and follow the wizard to commit your changes.
- ❹ Click on **CANCEL** to discard any modifications made in the fields and to select another entry in the configuration list or to add a whole new set of data.
- ❺ The list of existing configurations. The blue color indicates the selected line. During the modification, it turns gray.

Autocompletion Fields in the Wizards

Some fields in the wizards provide an autocompletion option, either manual or automatic.

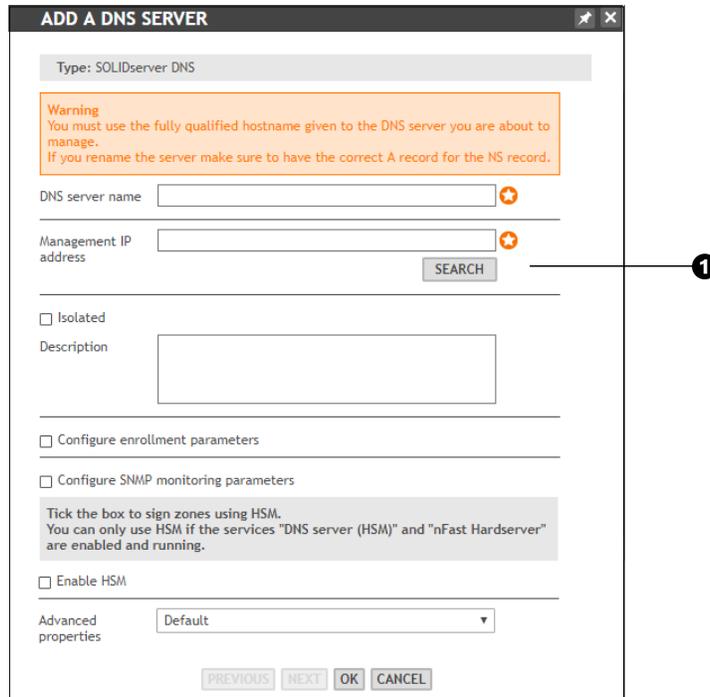


Figure 4.39. An example of manual autocompletion in the wizards

- 1 This field, above the button **SEARCH**, provides manual completion: in this example, you can type in a hostname to automatically retrieve its IP address if the DNS resolution is configured.

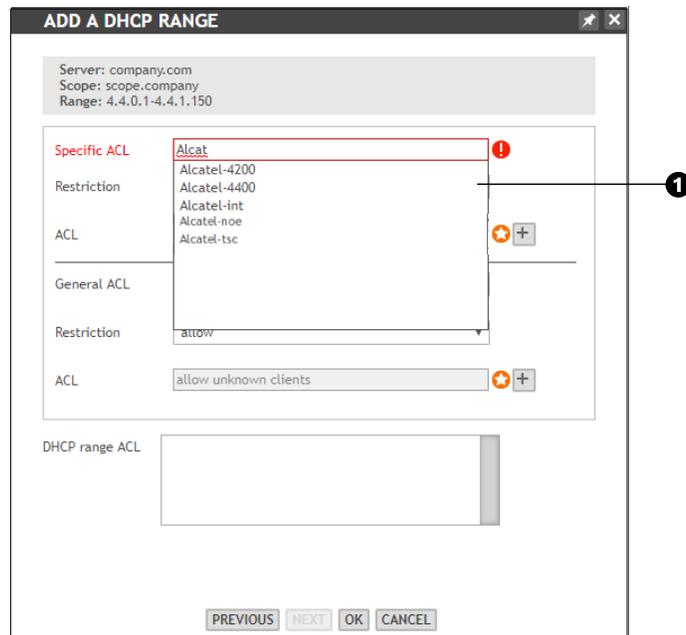


Figure 4.40. An example of automatic autocompletion in the wizards

- 1 There is no indicator of the autocompletion fields. Typing in a value automatically returns matching data in a list. If only one value matches your search, it is displayed in the field. If several entries match your search, a list appears under the field.

Quick Wizards

The quick wizard are shortcuts in essence that allow to save any wizard at any point of its configuration. The wizard's page and all the data filled or selected is saved within the quick wizard and accessible at any time.

The quick wizards are saved and managed on a dedicated page but you can set shortcuts toward each of them in a gadget or in a dedicated menu.

Browsing Quick Wizards

All quick wizards are saved and listed on the page *My Quick Wizards*, in the module *Administration*.

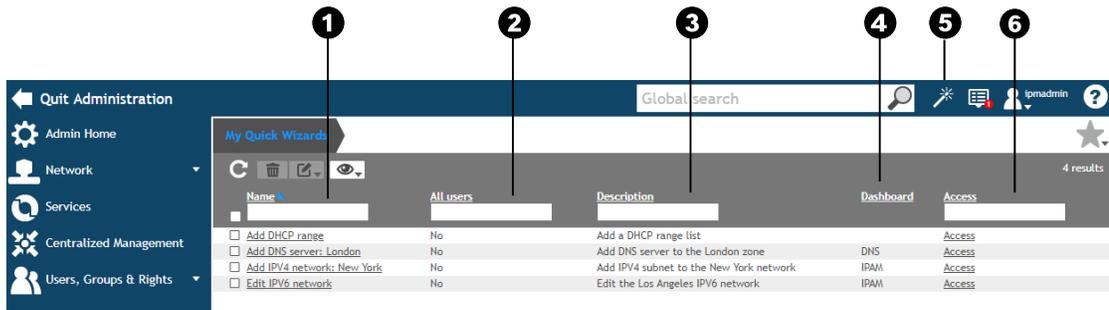


Figure 4.41. The page *My Quick Wizards*

- ❶ The column **Name** displays the quick wizard name. It allows to edit the quick wizards. For more details, refer to the section [Editing Quick Wizards](#).
- ❷ The column **All users** indicates if the quick wizard is shared with other users (*yes*) or not (*no*).
- ❸ The column **Description** displays the description you might have set during the quick wizard creation or edition.
- ❹ The column **Dashboard** indicates on which dashboard the Quick wizard gadget is displayed. It is empty if you only saved it in the *Quick access menu*. You cannot filter this column.
- ❺ The menu  **Quick access**. It appears if at least one quick wizard was assigned to the *Quick access menu*.
- ❻ The column **Access** is a link toward the wizard you saved.

The menu *Quick access*, can contain as many quick wizards as you want. You can access them from any page as the menu is displayed in the top bar.

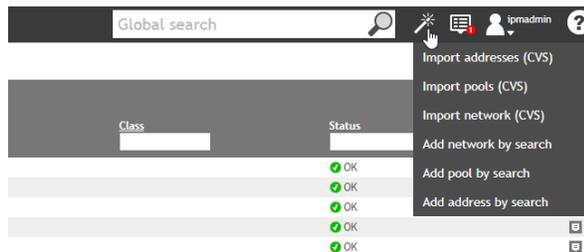


Figure 4.42. An example of *Quick Access menu*

To access the page My Quick Wizards

1. From any page, in the top bar, select **👤 My account > My Quick Wizards**. The page **My Quick Wizards** opens.
2. The page contains five columns that you can sort and/or filter. Quick wizards do not have a properties page, all the information is divided among the columns.

Adding Quick Wizards

The pushpin  located on the drag bar of all the pages of any wizard allows to create a quick wizard. Note that:

- When adding a quick wizard, you can only indicate one location for the shortcut, a module dashboard or the *quick access menu*. For more details regarding the gadget *Quick Wizards*, refer to the section [Creating a Quick Wizards Gadget](#).

You can edit it afterward to specify more locations. For more details, refer to the section [Editing Quick Wizards](#).

- Once created, a quick wizard is only available to the user who created it. To make it visible to everyone, you can edit its visibility. For more details, refer to the section [Editing Quick Wizards](#).
- You can access your quick wizards in different ways. For more details, refer to the section [Accessing Quick Wizards](#).

To add a quick wizard

1. From any wizard page, click on . The page **Add a Quick Wizard** opens.
2. In the field **Name**, type in the the quick wizard name.
3. In the drop-down list **Save in**, select either:
 - A module to display the gadget *Quick wizard* on the dashboard of the selected module. It includes a shortcut toward the quick wizard.
 - The menu *Quick access* to create a shortcut toward the quick wizard in the top bar .
4. In the field **Description**, you can type in a description.
5. Click on **OK** to complete the operation. The report opens, the wizard closes and the page you were on reopens. The quick wizard is listed on the page *My Quick Wizards* and is also accessible where you saved it.

Accessing Quick Wizards

There are three ways of accessing the quick wizards:

- Via the page *My Quick Wizards*.
- Via the gadget *Quick Wizards*. For more details, refer to the chapter [Managing Gadgets](#).
- Via the menu  *Quick Access* that is visible on every page of the appliance in the top bar.

To access a quick wizard from the page My Quick Wizards

1. From any page, in the top bar, select **👤 My account > My Quick Wizards**. The page **My Quick Wizards** opens.

2. Filter the list if need be.
3. In the column **Access**, click on *Access*. The wizard page you saved opens and all the data you filled before saving the wizard is already entered.

To access a quick wizard from the gadget **Quick Wizards**

1. Go to the dashboard where you displayed the gadget *Quick Wizards*.
2. In the gadget, click on the name of the quick wizard of your choice. The wizard page you saved opens and all the data you filled before saving the wizard is already entered.

To access a quick wizard from the menu **Quick Access**

1. Go to the listing page of your choice.
2. In the top bar, select  **Quick Access** > **<your-quick-wizard>**. The wizard page you saved opens and all the data you filled before saving the wizard is already entered.

Editing Quick Wizards

You can edit your quick wizards: you can rename them, change their description, their access method and visibility from the page *My Quick Wizards*.

To edit a quick wizard

1. From any page, in the top bar, select  **My account** > **My Quick Wizards**. The page **My Quick Wizards** opens.
2. Click on the name of the quick wizard you want to edit. The wizard **Quick Wizard Edition** opens.
3. If need be, edit the fields **Name** and **Description**.
4. In the list **Available**, select one by one the the locations of your choice.
5. Click on . Your selection is moved to the list **Configured**.
6. If you want to remove an item from the list **Configured**, select it and click on . The item is moved back to the list **Available**.
7. Tick or untick the box **Share with other users** according to your needs. Sharing a quick wizard makes it available for all users.
8. Click on to complete the operation. The report opens and closes.

Sharing Quick Wizards

You can share several quick wizards at once from the page *My Quick Wizards*.

To share several quick wizards at once

1. From any page, in the top bar, select  **My account** > **My Quick Wizards**. The page **My Quick Wizards** opens.
2. Tick the quick wizards you want to share.
3. In the menu, select  **Edit** > **Visible to all users** > **Yes**. The wizard **Quick wizard visibility** opens.
4. Click on to complete the operation. The report opens and closes.

To make several quick wizards visible only to you

1. From any page, in the top bar, select **My account > My Quick Wizards**. The page **My Quick Wizards** opens.
2. Tick the quick wizards you no longer want to share.
3. In the menu, select **Edit > Visible to all users > No**. The wizard **Quick wizard visibility** opens.
4. Click on **OK** to complete the operation. The report opens and closes.

Deleting Quick Wizards

You can only delete your quick wizards from the page *My Quick Wizards*.

To delete a quick wizard

1. From any page, in the top bar, select **My account > My Quick Wizards**. The page **My Quick Wizards** opens.
2. Tick the quick wizard(s) you want to delete.
3. In the menu, click on **Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The selected quick wizards are no longer listed.

Bookmarks

You can bookmark any page, even if it is displaying filtered data, and then manage them from their dedicated page or gadget.

Browsing Bookmarks

All bookmarks are listed and managed from the page *My Bookmarks*.

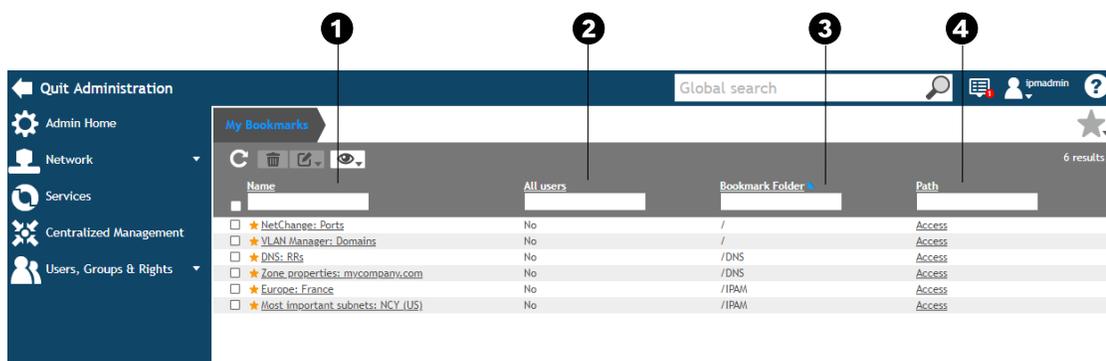


Figure 4.43. The page *My Bookmarks*

- 1 The column **Name** displays the bookmark name. It allows to edit the bookmarks. For more details, refer to the section [Editing Bookmarks](#).
- 2 The column **All users** indicates if you share the bookmark visibility with other users (Yes) or not (No).

- 3 The column **Bookmark Folder** indicates if the bookmark belongs to a folder. / means the bookmark is not in any folder.
- 4 The column **Path** contains the link *Access*, toward the bookmarked page.

To display the list of bookmarks

1. From any page, in the top bar, select **My account > My Bookmarks**. The page **My Bookmarks** opens.
2. Bookmarks do not have a properties page, all the information is divided among four columns that allow you to sort and/or filter the list.

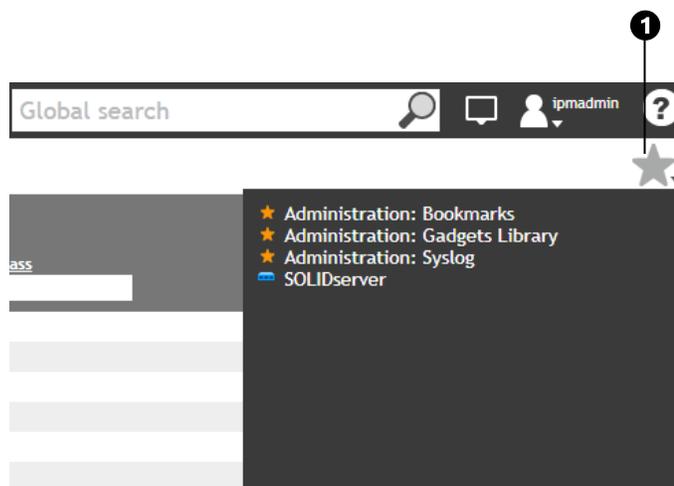


Figure 4.44. The window *Bookmarks*

- 1 When bookmarks are available in SOLIDserver, on the right-end side of the Breadcrumb, next to the icon ☆, a down arrow is displayed. It allows to display the list of all the available bookmarks.

To display the list of bookmarks from the window *Bookmarks*

1. From any page, on the right-end side of the breadcrumb, next to the icon ☆, click on the down arrow.
2. The list of bookmarks is displayed.

Adding Bookmarks

You can bookmark any page, you can even bookmark listing pages displaying filtered data.

To bookmark a page

1. From any page within SOLIDserver, on the right-end side of the breadcrumb, click on ☆. The wizard **Bookmark this page** opens.
2. In the field **Name**, specify your own bookmark name if need be. By default a bookmark is named *Module: Page*.
3. In the field **Bookmark Folder**, you can type in a folder name to organize the bookmarks on the page *My Bookmarks* and in the window *Bookmarks*.

The field allows to create or find folders: type in a folder name and click on to display the list of existing folders and select the one you need.

4. If you want to add the bookmark to the gadget Bookmarks, tick the box **Add to the gadget Bookmarks**.

For more details, refer to the sections [The Gadget Bookmarks](#) and [Creating a Gadget Bookmarks](#).

5. If you want to make the bookmark visible to any user, tick the box **Share with the other users**. If you leave it unticked, you are the only user who can see it.
6. Click on to complete the operation. The report opens and closes. The page is visible again and marked ★. The bookmark is listed on the page *My Bookmarks*.

Accessing Bookmarked Pages

Once you have bookmarked a page, you can access it from the page *My Bookmarks* or from the window *Bookmarks* in the Breadcrumb.

To access a bookmarked page from the page *My Bookmarks*

1. From any page, in the top bar, select **My account** > **My Bookmarks**. The page **My Bookmarks** opens.
2. At the end of the line of the bookmark of your choice, click on **Access**. The bookmarked page opens.

To access a bookmarked page from the window *Bookmarks*

1. From any listing or properties page, in the breadcrumb, click on the arrow next to the icon ☆. The list of all available bookmarks is displayed.
2. Click on the name of the bookmark of your choice. The corresponding page opens.

Editing Bookmarks

You can edit the bookmarks: rename them, place them in a (different) folder, attach them to the bookmark gadget and/or change the visibility settings from the page *My Bookmarks*.

To edit a bookmark details

1. From any page, in the top bar, select **My account** > **My Bookmarks**. The page **My Bookmarks** opens.
2. Click on the name of the bookmark you want to edit. The wizard **Edit Bookmarks** opens.
3. If need be, edit the fields **Name** and **Bookmark Folder**.
4. Tick/untick the boxes **Add to the gadget Bookmarks** (if this box is ticked the gadget Bookmarks includes a shortcut toward the page) and **Share with the other users** (if the box is ticked all the users can see your bookmark).
5. Click on to complete the operation. The report opens and closes. The content of the columns matches your modifications.

Sharing Bookmarks

From the page *My Bookmarks* you can share one or several bookmarks at once.

To share bookmarks

1. From any page, in the top bar, select  **My account** > **My Bookmarks**. The page **My Bookmarks** opens.
2. Tick the bookmark(s) of your choice.
3. In the menu, select  **Edit** > **Visible to all users** > **Yes**. The wizard **Bookmark Visibility** opens.
4. Click on to complete the operation. The report opens and closes. The page is visible again. The bookmark is marked *Yes* in the column *All users*.

To make bookmarks visible only to you

1. From any page, in the top bar, select  **My account** > **My Bookmarks**. The page **My Bookmarks** opens.
2. Tick the bookmark(s) of your choice.
3. In the menu, select  **Edit** > **Visible to all users** > **No**. The wizard **Bookmark Visibility** opens.
4. Click on to complete the operation. The report opens and closes. The page is visible again. The bookmark is marked *No* in the column *All users*.

Deleting Bookmarks

You can delete bookmarks from the bookmarked page or from the page *My Bookmarks*. Note that you can delete several bookmarks at one from this page.

To delete a bookmark from the page itself

1. Go to the bookmarked page,
2. On the right-end side of the breadcrumb, click on . The wizard **Delete** opens.
3. Click on to complete the operation. The page refreshes. The page is marked .

To delete one or more bookmarks from the page My Bookmarks

1. From any page, in the top bar, select  **My account** > **My Bookmarks**. The page **My Bookmarks** opens.
2. Tick the bookmark(s) you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on to complete the operation. The report opens and closes. The page is visible again.

Part II. Configuring SOLIDserver

Before managing your network, your administrator needs to configure your appliance.

This part details all the available system configurations needed to set up SOLIDserver from the module Administration, they are divided as follows:

- [Configuring the Time and Date](#): describes the ways of setting the appliance time and date, a mandatory configuration to ensure services synchronization and data reliability.
- [Configuring the Network](#): describes the operations to integrate the appliance to your network. From its IP address and hostname to its DNS resolver, firewall settings, routes and so on.
- [Configuring the Services](#): describes all the services and servers that you can configure and/or manage from SOLIDserver: SSH, NTP, HTTP, DNS, DHCP, etc.
- [Managing the Licenses](#): describes how to renew or delete license(s).

Note that the module Administration provides extra pages and features all described in the parts [Administration](#) and [Customization](#).

Chapter 5. Configuring the Time and Date

Your appliance must always be set with the proper time and date to prevent any management problems. That way, all your services are properly synchronized and all the data you manage is up-to-date.

There are two ways of configuring the appliance time and date:

1. Via NTP.

We strongly recommend configuring NTP servers on your appliance. You can configure several servers and even force an update. For more details, refer to the sections [Configuring NTP Servers](#) and [Forcing the NTP Update](#).

2. Manually.

You can set the date and time yourself as detailed in the section [Setting the Appliance Time and Date Manually](#).

Note that every user can choose time and date display of their session. For more details, refer to the section [Configuring the User Display Settings](#).

Configuring NTP Servers

The Network Time Protocol (NTP) ensures clock synchronization on a network. You must configure NTP servers on your appliance to make sure that its services are set with the proper time and date, thus ensuring that any transfer, exchange or synchronization of DHCP, DNS, SNMP or high availability data is possible. In other words, NTP servers ensure that all the data managed from SOLIDserver is reliable and up-to-date.

When configuring the NTP server, keep in mind that:

- You can configure your appliance with public or private NTP servers and each server can have a specific stratum level.
- You should configure at least 3 reference NTP servers for all the NTP clients on your network.
- The reference NTP servers must be reachable when you start the service.
- All the services must be set at the same time to prevent any management problems.
- You can force an NTP update at any time, for more details refer to the section [Forcing the NTP Update](#).

Note that you can configure NTP servers from the page *Network configuration* in the module Administration, as detailed below, or from the gadget *SOLIDserver Configuration Checklist* on the *Main Dashboard*.

To configure NTP servers

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.

3. Under the menu, in the drop-down list **SOLIDserver**, make sure the appliance of your choice is selected.
4. In the column **Name**, click on the link **NTP server**. The wizard **NTP Servers Configuration** opens.
5. In the field **NTP address**, type in the IP address or hostname of the server. It can be an IPv4 or IPv6 address.
6. In the field **Stratum**, you can specify a level between 0 and 15. By default nothing is specified, the stratum is retrieved from the server. We strongly advise against setting a stratum if it is not necessary.
7. Click on **ADD** to move the data in the list **NTP servers**.
 - To update an entry, select the NTP server of your choice, change data and click on **UPDATE**.
 - To delete an entry, select the NTP server of your choice and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
8. Repeat these steps to add as many servers as you need.
9. With at least two servers in the list **NTP servers**, you can set in which order the servers are interrogated. Select one and move it up or down the list.
10. Click on **OK** to complete the operation.
11. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

If you need to edit the NTP servers configuration, follow the procedure again and make your changes. To take into account your changes immediately, you can stop the service NTP and start it again. For more details refer to the section [Starting or Stopping a Service](#).

Forcing the NTP Update

At any time, you might need to force an update of the NTP servers time and date.

Before forcing an NTP update:

- **Make sure that at least one NTP server is configured and reachable**, otherwise you might not be able to access your appliance at all.
- **Keep in mind that forcing the update restarts all the services that rely on NTP**, like the services DNS, DHCP and SNMP.

To force an NTP update

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. Under the **NTP server** line, click on **FORCE UPDATE**. The wizard **Force NTP update** opens.
4. Click on **OK** to complete the operation. The services restart.

Setting the Appliance Time and Date Manually

From the Administration homepage, you can set the time and date of your appliance without relying on NTP.

We recommend configuring NTP servers on your appliance to make sure that the time and date are regularly checked and updated, for more details, refer to the section [Configuring NTP Servers](#).

Note that, if you set up one or several NTP servers, the time and date you set manually will be lost the next time an NTP server updates.

To manually set the appliance time and date

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Time & Date configuration**. The wizard opens.
3. Set the appliance time and date, by default each field is prefilled with a value matching the current time and date of the appliance:
 - a. In the drop-down list **Hour**, select a value between 0 and 24.
 - b. In the drop-down list **Minute**, select a value between 0 and 59.
 - c. In the drop-down list **Day**, select the date of your choice, a value between 0 and 31.
 - d. In the drop-down list **Month**, select the value of your choice.
 - e. In the field **Year**, type in the value of your choice.
4. Click on **OK** to complete the operation. The services restart.

Chapter 6. Configuring the Network

This chapter details the page **Network configuration** where you can configure all the settings necessary to run SOLIDserver on your network, including:

- [Setting the Hostname](#) of the appliance.
- [Setting the DNS Resolver](#), the DNS server that SOLIDserver uses to resolve the names and addresses that it manages.
- [Setting the Firewall](#) and reinforcing the appliance security by blocking potential dangerous communications.
- [Setting up the Default Gateway](#) address that SOLIDserver uses to reach networks out of its domain's broadcast.
- [Setting up Specific Routes](#) to set a specific path for the returned packets.
- [Setting up Static Routes](#) and enable data to be forwarded through the network with fixed paths.
- [Configuring Basic IP Addressing on an Interface](#).
- [Setting up a VLAN Interface](#)¹, like using a physical interface as an 801.1Q interface.
- [Setting up an Ethernet Port Failover](#), to allow aggregation of multiple network interfaces as one virtual interface in order to provide fault-tolerance and high-speed links.
- [Configuring a VIP](#), or Virtual IP, that is not connected to a specific computer or network interface card on a computer. Incoming packets are sent to the VIP address, but all packets travel through real network interfaces.
- [Setting up a VIF](#), or Virtual InterFace, an EfficientIP feature that allows to add into a VIF a configuration of physical interfaces embedding many services.
- [Configuring a Media Interface](#), to define the option supported by the physical interface.

Keep in mind that:

- **All the services must be set at the same time to prevent any management problems.** For more details, refer to the section [Configuring NTP Servers](#).
- **You can set the network of remote appliances** via the drop-down list *SOLIDserver*, whether you simply manage an appliance remotely or configured it in high availability. For more details, refer to the chapter [Centralized Management](#).

¹Virtual Local Area Network (VLAN) is a group of hosts with a common set of requirements that communicate as if they were attached to the same broadcast domain, regardless of their physical location.

Configuring Basic IP Addressing on an Interface

Multiple IP addresses can be configured on a single VIF. Configuring multiple IP addresses on a VIF can be helpful in different scenarios, such as DNS server migration. Configuring multiple IP addresses of existing DNS servers to a single VIF enables administrators to provide continuous service during server migration or High Availability of the service through different appliances. By default, an existing VIF (called `DEFAULT_INTERFACE`) is already applied in the system, you can use this one or create a new one. In order to apply a new one, refer to the section [Setting up a VIF](#).

Keep in mind that **the overlap of IP addresses linked with different physical interfaces is not allowed** in order to avoid asymmetrical routing. Indeed, if a packet is received from a physical interface it must not be forwarded to another one.

To set up a Basic Interface Configuration

Only users of the group `admin` can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the link `DEFAULT_INTERFACE`. The wizard **Virtual network interface configuration** opens.
4. In the field **Virtual interface name**, you can rename the default interface if you want.
5. In the list **Available physical interfaces**, select the available interface, it is named after the physical port and port MAC address as follows: `eth# (##:##:##:##:##:##)` and click on . It is now listed in the field **Physical interfaces**.
6. Click on `NEXT`. The next page opens.

If you selected at least two *Physical interfaces*, in the drop-down list **LAGG protocol** you can select *failover* or *LACP*. By default, *failover* is selected. Click on `NEXT`.

Note that a successful LAGG configuration requires interfaces with the same speed and duplex and you can only configure LACP on appliances in version 6.0.2 or higher.

7. Fill in the interface IPv4 address configuration parameters following the table below:

Table 6.1. IPv4 virtual network interface configuration parameters

Parameter	Description
IP address	Type in the interface IP address. This field is mandatory.
Netmask	Type in the interface netmask. This field is mandatory.
Specific route	This field allows Setting up Specific Routes (setting up source routing) if necessary. The specified route is dedicated to the IP address.

Once all the parameters needed are configured, click on `ADD`. The new IP address is now listed in the field **IP addresses list**. You can add multiple IP addresses. SOLIDserver can be accessed through all the IP addresses configured on this VIF.

- To update an entry, select a configured IP address, change the needed data and click on `UPDATE`.
- To delete an entry, select a configured IP address and click on `DELETE`.

- To discard the latest modifications, click on **CANCEL**.
8. Click on **NEXT**. The last page of the wizard opens. Fill in the interface IPv6 address configuration parameters following the table below:

Table 6.2. IPv6 virtual network interface configuration parameters

Parameter	Description
IPv6 address	Type in the interface IP address. This field is mandatory.
Prefix	Type in the interface prefix. This field is mandatory.

Once all the parameters needed are configured, click on **ADD**. The new IP address is moved to the **IPv6 addresses list**. You can add multiple IP addresses. SOLIDserver is accessible through all the IP addresses configured on this VIF.

9. Click on **OK** to complete the operation. If you configured LAGG, the protocol you chose is displayed in the column **Configuration**.
10. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them.

Make sure that at least one interface is available, otherwise, you would lose your current connection to SOLIDserver.

Setting the Routing

With SOLIDserver you can configure routing on your network by:

- [Setting up the Default Gateway](#), to forward SOLIDserver outgoing traffic.
- [Setting up Specific Routes](#), to set a path for returned packets.
- [Setting up Static Routes](#), to forward data to another subnet.

Setting up the Default Gateway

A gateway is a node on a TCP/IP network that is used as an access point to another network. The default router is the gateway used by SOLIDserver that forwards traffic to remote subnets on behalf of a sending host or router. Only one default router can be configured for the entire appliance in each version of the IP protocol. For security reasons, SOLIDserver does not route packets between network interfaces.

Keep in mind that the default gateway is only used if a packet is sent from a network address unknown to SOLIDserver. For some networks, you might want to use route sourcing and set up a specific route to send the response packet to the sender through the channel it came from rather than using the default gateway to try and locate the sender. For more details, refer to the section [Setting up Specific Routes](#).

To configure the default gateway

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.

3. Click on the link **Default gateways** in the network configuration listing. The wizard **Edit the default gateways** opens.
4. In the field **IPv4 default gateway**, fill in the IPv4 gateway of your choice.
5. In the field **IPv6 default gateway**, fill in the IPv6 gateway of your choice.
6. Click on to complete the operation.
7. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Setting up Specific Routes

Setting up specific routes allows to configure source routing on your network for IPv4. Without specific routes, once a packet is sent from a subnet that has not been configured among SOLIDserver network interfaces, the response packet is returned through the default gateway and might, depending on your network architecture, never get back to the sender (asymmetric routing issue). In other words, setting up a specific route allows to specify the route for the return packet. Once a specific route is configured, its address is preferred to the default gateway for the return packets.

Keep in mind that the DHCP does not take into account the specific route. Therefore, the management IP address a DHCP server should always be on the same network as the default gateway.

Within SOLIDserver, you can set up several specific routes. To configure a specific route, refer to the procedure [To set up a Basic Interface Configuration](#).

Setting up Static Routes

If it is necessary, SOLIDserver allows you to add static routes. These routes allow you to communicate with another network(s) and to forward data through a fixed path.

To configure static routes (Add/Edit/Delete)

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the link **Static routes**. The page **Static Routes (IPv4)** opens.

Table 6.3. IPv4 static route configuration parameters

Parameter	Description
Route name	Name the static route. This field is mandatory.
IP address	Type in the static route IP address. This field is mandatory.
Netmask	Depending on the IP address you typed in above, you might have a list of netmasks to choose from. The netmask you choose automatically selects the corresponding prefix. This field is mandatory.
Prefix	Depending on the IP address and selected netmask, a prefix is automatically selected. If you choose a different prefix, the netmask is modified accordingly. This field is optional.
Gateway	Type in the gateway you want to use with the static route. This field is mandatory.

Once all the parameters needed are configured, click on **ADD**. The static route is now listed in the list **Static routes**. You can add multiple static routes.

- To update an entry, select an existing static route, change data and click on **UPDATE**.
 - To delete an entry, select an existing static route and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
4. Click on **NEXT**. The page **Static routes (IPv6)** opens. Follow the step 4 to configure an IPv6 static route.
 5. Click on **OK** to complete the operation.
 6. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

Setting the Hostname

The hostname in the name of you local appliance. It must be a Fully Qualified Domain Name (FQDN), in other words the name of the host concatenated with the domain name.

The hostname is used to identify and differentiate several appliances. It is all the more useful if you manage remote appliances or configure appliances in high availability.

To configure an appliance hostname

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the link **Hostname** in the network configuration listing. The wizard **Edit the hostname** opens.
4. In the field **Hostname**, name your hostname with a valid FQDN. By default, every appliance is named *solid.intranet*.
5. Click on **OK** to complete the operation.
6. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

Setting the DNS Resolver

The DNS resolver is the default DNS server that SOLIDserver uses to resolve local names. Several modules like IPAM, NetChange and DNS Manager use the DNS resolver to find IP addresses' FQDN or to resolve an FQDN IP address.

To configure DNS resolvers (Add/Edit/Delete)

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the link **DNS Resolvers**. The wizard **Edit DNS resolvers** opens.
4. In the field **DNS server IP address**, type in the IP address of the server of your choice and click on **ADD**.
5. Click on **ADD**. The IP address is now moved to the list **DNS Resolvers**.
6. Repeat these operations for as many resolvers as needed. If you have several resolvers use the **▲** and **▼** buttons to order the list according to your needs.
 - To update an entry, select a DNS resolver, change the needed data and click on **UPDATE**.
 - To delete an entry, select a DNS resolver and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
7. Click on **OK** to complete the operation.
8. Right now your configuration is pending. In the menu, select **⚙️ Tools > Apply configuration** to save your changes or **⚙️ Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

Setting the Firewall

SOLIDserver embeds a restrictive *stateful firewall*² for securing flows. SOLIDserver firewall uses the legacy stateless rules and a legacy rule coding technique to achieve what is referred to as *Simple Stateful* logic.

SOLIDserver stateful filtering treats traffic as a bi-directional exchange of packets comprising a session conversation. It has the matching capabilities to determine if the session conversation between the originating sender and the destination are following the valid procedure of bi-directional packet exchange. Any packets that do not properly fit the session conversation template are automatically rejected. SOLIDserver allows firewall messages filing making it possible to review after the fact information such as: which packets have been dropped, from which addresses they came from and where they were going, giving you significant capacity to track down attackers. SOLIDserver supports *Stateful Packet Inspection* (SPI) mode that helps preventing network attacks by tracking more state per session.

Enabling or Disabling the Firewall

By default, SOLIDserver firewall is *Restricted*, i.e. enabled, and all the firewall rules set are respected and enforced in order.

At any time, you can *Open* the firewall, to disable it, and ignore all these rules.

To open or restrict the firewall

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **⚙️ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.

²State full Packet Inspection, also known as dynamic packet filtering.

3. In the column **Configuration**, in the line **Firewall** can be *Restricted* or *Open*.
4. Click on the current state to change it. The wizard **Firewall state configuration** opens.
5. Click on to complete the operation. The firewall is marked *Open* or *Restricted*.
6. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Adding a Firewall Rule

Before adding or editing firewall rules, you need to understand a few concepts:

Precedence

It is a key concept in the firewall rule configuration. It corresponds to a number specified in the Firewall rule wizard. All the parameters that you configure (action, protocol, from, to, port, via, log and keep state) in the wizard set up a distinct set of conditions that, if matched, is dealt with respecting the order set in the field Position. Therefore it is paramount to understand that if for instance you set two firewall rules regarding the ipv4/ipv6 protocol from a DNS server A to a DNS server B through the port 53 via em0, and one denies access whereas the other accepts it, the rule that prevails is the one set with the smallest position number of the two in the field Position.

Firewall rules

The firewall being restrictive, as opposed to permissive, the last position (65535) denies access to any kind of packets no matter what protocol or where it goes or comes from. Which is why EfficientIP has configured a number of firewall rules, they are all listed on the page *Firewall rules*. On this page you can edit a number of preexisting rules and the ones you create: the underlined rules in the column **Position** can be edited, all the others cannot. For technical reasons, the positions 1 - 99 are reserved by EfficientIP and users cannot use any of them when creating rules or editing rules. The position 65535 cannot be used either.

To add a firewall rule

Only users of the group *admin* can perform this operation.

1. Make sure you understand the [Precedence](#) and [Firewall rules](#).
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Firewall rules**. The page **Firewall rules** opens.
4. In the menu, click on  **Add**. The Firewall rule configuration wizard opens, fill in all the required parameters following the table below:

Table 6.4. Firewall rules parameters

Parameter	Description		
Position	Set the rule precedence using a number between 100 and 65534. For more details, refer to the paragraphs Precedence and Firewall rules above.		
Action	Define what action should be executed when a packet matches the selection criterion of the rule. For each rule you can: <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 2px;"><i>allow</i>: packets matching the defined criterion. The rule exits the firewall rule processing. The search terminates at this rule.</td> </tr> <tr> <td style="padding: 2px;"><i>deny</i>: packets matching the defined criterion. The packets are discarded. The search terminates.</td> </tr> </table>	<i>allow</i> : packets matching the defined criterion. The rule exits the firewall rule processing. The search terminates at this rule.	<i>deny</i> : packets matching the defined criterion. The packets are discarded. The search terminates.
<i>allow</i> : packets matching the defined criterion. The rule exits the firewall rule processing. The search terminates at this rule.			
<i>deny</i> : packets matching the defined criterion. The packets are discarded. The search terminates.			

Parameter	Description
Protocol	In this drop-down list, choose the protocol used for that rule. These protocols will handle IPv4 and/or IPv6 protocols.
From	Define the source parameters. The fields <i>From</i> and <i>To</i> work together, so you must specify either two IPv4 addresses or two IPv6 addresses, you cannot mix the protocol versions. Accepted values are: <i>me</i> : a special keyword that matches any IP address configured on an interface in SOLIDserver. <i>any</i> : a special keyword that matches any IP address. <IP-address> specified with mask-length following the format: x.x.x.x/x or xxxx::/x . <IP-address> specified without mask-length following the format: x.x.x.x or xxxx:: .
Source port	Define the source port on which the firewall rule should be applied. Use a comma to separate several port numbers.
To	Define the destination parameters. The fields <i>From</i> and <i>To</i> work together, so you must specify either two IPv4 addresses or two IPv6 addresses, you cannot mix the protocol versions. Accepted values are: <i>me</i> : a special keyword that matches any IP address configured on an interface in SOLIDserver. <i>any</i> : a special keyword that matches any IP address. <IP-address> specified with mask-length following the format: x.x.x.x/x or xxxx::/x . <IP-address> specified without mask-length following the format: x.x.x.x or xxxx:: .
Destination port	Define the destination port on which the firewall rule should be applied. Use a comma to separate several port numbers.
Via	Set the interface the packets should go through. The via parameter causes the interface to always be checked as part of the match process. By default, nothing is selected.
Log	Choose to save, or not, the log parameter indicating if a packet matches a rule in SOLIDserver syslog page (it is saved with a facility SECURITY name). By default, No is selected.
Keep state	Decide if you want SOLIDserver firewall to create a dynamic rule, upon match, whose default behavior is to match bidirectional traffic between source and destination IP/port using the same protocol. By default, No is selected.

- Click on to complete the operation.
- Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Editing a Firewall Rule

You can edit the firewall rules you added. Before editing a rule, make sure you understand the [Precedence](#) and [Firewall rules](#) detailed in the section [Adding a Firewall Rule](#).

To edit a firewall rule

Only users of the group *admin* can perform this operation.

- Make sure you understand the [Precedence](#) and [Firewall rules](#).
- In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **System**, click on **Firewall rules**. The page **Firewall rules** opens.

4. In the column **Position**, click on the underlined number corresponding to the rule you want to edit. The wizard **Firewall rule configuration** opens.
5. Edit the parameters according to your needs, following the information described in [Firewall rules parameters](#) procedure above.
6. Click on to complete the operation.
7. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Deleting a Firewall Rule

Any firewall rule can be deleted, except the rule 65535.

Keep in mind that **firewall rules must not be deleted lightly**. For instance, the rule #34 is a delicate rule to delete as it refers to fragmented IP packets. As there is a maximum packet size for transport level that depends on the transport medium (1500 bytes for Ethernet), if the IP packet is larger than this, it needs to be broken up into fragments that get reassembled at the destination. Without the rule #34, fragmented IP packets will be blocked by the firewall.

To delete a firewall rule

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Firewall rules**. The page **Firewall rules** opens.
3. Tick the firewall rule you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation.
6. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Setting up a VLAN Interface

You can set up a VLAN interface via the VIF configuration wizard.

By default, a VIF (called *DEFAULT_INTERFACE*) is already created in the network configuration, you can edit it or create a new one. To do so, refer to the section [Setting up a VIF](#).

Note that **to avoid asymmetrical routing, you cannot link overlapped IP addresses to different physical interfaces**. This way, if a packet is received from a physical interface it cannot be forwarded to another interface.

To set up a VLAN interface configuration

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of the interface you want to configure (all interfaces are preceded by an orange dot). The wizard **Virtual network interface configuration** opens.
4. In the field **Virtual interface name**, you can rename the default interface if you want.
5. In the list **Available physical interfaces**, select the available interface, it is named after the physical port and port MAC address as follows: *eth# (##:##:##:##:##:##)* and click on **+**. It is now listed in the field **Physical interfaces**.
6. Click on **NEXT**. The next page opens.

If you selected at least two *Physical interfaces*, in the drop-down list **LAGG protocol** you can select *failover* or *LACP*. By default, *failover* is selected. Click on **NEXT**.

Note that a successful LAGG configuration requires interfaces with the same speed and duplex and you can only configure LACP on appliances in version 6.0.2 or higher.

7. Fill in the interface IPv4 addresses configuration parameters following the table below:

Table 6.5. IPv4 virtual network interface configuration parameters

Parameter	Description
IP address	Type in the interface IP address. It should correspond to one of the VLAN configured on your network. This field is mandatory.
Netmask	Type in the interface netmask. This field is mandatory.
Specific route	This field allows Setting up Specific Routes (setting up source routing) if necessary. The specified route is dedicated to the IP address.
802.1q tag number	Type in the VLAN number of your choice (between 1 and 4094). This tag can be common to different appliances and differentiate them from other IP addresses on the VLAN: packets sent to the VLAN with the same tag are only received by these appliances.

Once all the parameters needed are configured, click on **ADD**. The new IP address is now listed in the field **IP addresses list**. You can add multiple IP addresses. The IP determines to which configured VLAN they belong and the tag provides a more accurate filter.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
- To delete an entry, select a configured IP address and click on **DELETE**.
- To discard the latest modifications, click on **CANCEL**.

8. Click on **NEXT**. The last page of the wizard opens. Fill in the interface IPv6 address configuration parameters following the table below:

Table 6.6. IPv6 virtual network interface configuration parameters

Parameter	Description
IPv6 address	Type in the interface IP address. It should correspond to one of the VLAN configured on your network. This field is mandatory.
Prefix	Type in the interface prefix. This field is mandatory.
Specific route	You can apply a Specific route (set up source routing) if necessary. This route is dedicated to the IP address.
802.1q tag number	Type in the VLAN number of your choice (between 1 and 4094). This tag can be common to different appliances and differentiates them from other IP addresses on

Parameter	Description
	the VLAN: packets sent to the VLAN with the same tag are only received by these appliances.

Once all the parameters needed are configured, click on **ADD**. The new IP address is moved to the **IPv6 addresses list**. You can add multiple IP addresses, the IP determines to which configured VLAN they belong and the tag provides a more accurate filter.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
 - To delete an entry, select a configured IP address and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
9. Click on **OK** to complete the operation. If you configured LAGG, the protocol you chose is displayed in the column **Configuration**.
 10. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them.

Make sure that at least one interface is available, otherwise, you would lose your current connection to SOLIDserver.

Setting up an Ethernet Port Failover

The Ethernet Port Failover is an ability of the network system to have 2 or more physical interfaces configured with one (or more) IP address access. To sum up, Ethernet Port Failover interface ensures a high SOLIDserver accessibility (if one of the physical interfaces is disconnected, the system is still available). By default, an existing VIF (called *DEFAULT_INTERFACE*) is already applied in the system, you can use this one and create others to set up a failover. For more details regarding interface addition, refer to the section [Setting up a VIF](#).

To configure an Ethernet Port Failover Interface Configuration

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of the interface you want to configure (all interfaces are preceded by an orange dot). The wizard **Virtual network interface configuration** opens.
4. In the field **Virtual interface name**, you can rename the default interface if you want.
5. In the list **Available physical interfaces**, select two or more interfaces one by one and click on **+**. They are now listed in the field **Physical interfaces**.
6. Click on **NEXT**. The next page opens.

If you selected at least two *Physical interfaces*, in the drop-down list **LAGG protocol** you can select *failover* or *LACP*. By default, *failover* is selected. Click on **NEXT**.

Note that a successful LAGG configuration requires interfaces with the same speed and duplex and you can only configure LACP on appliances in version 6.0.2 or higher.

7. Fill in the interface IPv4 address configuration parameters following the table below:

Table 6.7. IPv4 virtual network interface configuration parameters

Parameter	Description
IP address	Type in the interface IP address. This field is mandatory.
Netmask	Type in the interface netmask. This field is mandatory.
Specific route	This field allows Setting up Specific Routes (setting up source routing) if necessary. The specified route is dedicated to the IP address.

Once all the parameters needed are configured, click on **ADD**. The new IP address is now listed in the field **IP addresses list**.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
 - To delete an entry, select a configured IP address and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
8. Click on **NEXT**. The last page of the wizard opens. Fill in the interface IPv6 address configuration parameters following the table below:

Table 6.8. IPv6 virtual network interface configuration parameters

Parameter	Description
IPv6 address	Type in the interface IP address. This field is mandatory.
Prefix	Type in the interface prefix. This field is mandatory.
Specific route	You can apply a specific route (set up source routing) if necessary. This route is dedicated to the IP address.

Once all the parameters needed are configured, click on **ADD**. The new IP address is moved to the **IPv6 addresses list**.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
 - To delete an entry, select a configured IP address and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
9. Click on **OK** to complete the operation. If you configured LAGG, the protocol you chose is displayed in the column **Configuration**.
10. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them.

Make sure that at least one interface is available, otherwise, you would lose your current connection to SOLIDserver.

Configuring a VIP

By default, an existing VIF, called *DEFAULT_INTERFACE*, is already applied in the system, you can use this one or create a new one. In order to apply a new one, refer to the section [Setting up a VIF](#).

SOLIDserver allows you to set up virtual IP addresses (VIP) on supported services. This mechanism, known as *Common Address Redundancy Protocol* (CARP) is a protocol which allows multiple EfficientIP devices on the same local network to share a single IP address or the same

set of addresses. Its primary purpose is to provide failover redundancy. For example, if there is a single SOLIDserver running a DNS service and it goes down, then, the networks on each side of the DNS service can no longer communicate with each other, or, they communicate without any DNS service. However, if there are two EfficientIP devices running CARP, if one fails, the other can take over with SOLIDserver on either side of the DNS service not being aware of the failure. Operations continue as normal. Note that through a VIP you can manage DNS smart architectures *Master/Slave* and *Multi-Master*.

The general idea is to have a single IP address, and several physical servers behind. In the case of a failure, the next available server takes the lead and provides the relevant services. This mechanism is available for DNS, NTP, TFTP services and SOLIDserver management.

Please note that, when using virtual appliances, the VMware ESXi host vSwitch must be configured with the *Promiscuous mode* option *enabled* and the option *Net.ReversePathFwdCheckPromisc* set to *1*.

To configure a VIP

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of the interface you want to configure (all interfaces are preceded by an orange dot). The wizard **Virtual network interface configuration** opens.
4. In the field **Virtual interface name**, you can rename the default interface if you want.
5. In the list **Available physical interfaces**, select the available interface, it is named after the physical port and port MAC address as follows: *eth# (##.##.##.##.##.##)* and click on . It is now listed in the field **Physical interfaces**.
6. Click on **NEXT**. The next page opens.

If you selected at least two *Physical interfaces*, in the drop-down list **LAGG protocol** you can select *failover* or *LACP*. By default, *failover* is selected. Click on **NEXT**.

Note that a successful LAGG configuration requires interfaces with the same speed and duplex and you can only configure LACP on appliances in version 6.0.2 or higher.

7. Fill in the interface IPv4 address configuration parameters following the table below:

Table 6.9. IPv4 virtual network interface configuration parameters

Parameter	Description
IP address	Type in the interface IP address. This field is mandatory.
Netmask	Type in the interface netmask. This field is mandatory.
Specific route	This field allows Setting up Specific Routes (setting up source routing) if necessary. The specified route is dedicated to the IP address.
VIP service	You can select <i>DNS server</i> , <i>NTP server</i> , <i>TFTP server</i> or <i>SOLIDserver management</i> . By default, <i>None</i> is selected.
DNS server	The IP address is dedicated to the service you select. Selecting one of these services displays the fields <i>VHID</i> , <i>Password</i> and <i>Priority</i> .
NTP server	
TFTP server	

Parameter	Description
SOLIDserver management	Selecting this service allows you to access the appliances configured in High Availability through the IP address configured above. This virtual IP address only gives you access to the Master appliance. For more details, refer to the chapter Centralized Management . Selecting this service server displays the fields <i>VHID</i> and <i>Password</i> ^a .
VHID	Type in the Virtual Host IDentification if you are setting up the High Availability of the selected service. This VHID must be a number between 1 and 255 and it has to be the same on the appliances through which you set the service High Availability.
Password	Type in the password of your choice if you are setting up the High Availability of the selected service. This password has to be the same on the appliances set in high availability.
Priority	You can set the appliance priority to <i>Low</i> , <i>Medium</i> or <i>High</i> . In other words, you can decide which appliance has the priority over a service.

^aUnlike the other services, the Priority of SOLIDserver management cannot be modified: the Master always has priority over the Hot Standby.

These fields allow to set up the availability of the DNS, NTP, TFTP or SOLIDserver management services as long as both appliances belong to the same LAN (layer 2) and both appliances are set with the exact same parameters in all the fields EXCEPT for the *Priority*. To avoid any conflict, you must set one priority level for the first appliance and a different one on the other.

Once all the parameters needed are configured, click on **ADD**. The new IP address is now listed in the field **IP addresses list**. You can add multiple IP addresses. SOLIDserver is accessible through all the IP addresses configured on this VIF.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
- To delete an entry, select a configured IP address and click on **DELETE**.
- To discard the latest modifications, click on **CANCEL**.

8. Click on **NEXT**. The last page of the wizard opens. Fill in the interface IPv6 address configuration parameters following the table below:

Table 6.10. IPv6 virtual network interface configuration parameters

Parameter	Description
IPv6 address	Type in the interface IP address. This field is mandatory.
Prefix	Type in the interface prefix. This field is mandatory.
Specific route	You can apply a specific route (set up source routing) if necessary. This route is dedicated to the IP address.
VIP service	You can select one of four services: the <i>DNS server</i> , <i>NTP server</i> or <i>TFTP server</i> . By default, <i>None</i> is selected.
DNS server	The IP address is dedicated to the service you select. Selecting one of these services displays the fields <i>VHID</i> , <i>Password</i> and <i>Priority</i> .
NTP server	
TFTP server	
VHID	Type in the Virtual Host IDentification if you are setting up the High Availability of the selected service. This VHID must be a number between 1 and 255 and it has to be the same on the appliances through which you set the service High Availability.
Password	Type in the password of your choice if you are setting up the High Availability of the selected service. This password has to be the same on the appliances set in High Availability.
Priority	You can set the appliance priority to <i>Low</i> , <i>Medium</i> or <i>High</i> . In other words, you can decide which appliance has the priority over a service.

These fields allow to set up the availability of the DNS, NTP, TFTP services as long as both appliances belong to the same LAN (layer 2) and both appliances are set with the exact same parameters in all the fields EXCEPT for the *Priority*. To avoid any conflict, you must set one priority level for the first appliance and a different one on the other.

Once all the parameters needed are configured, click on **ADD**. The new IP address is moved to the **IPv6 addresses list**. You can add multiple IP addresses. SOLIDserver is accessible through all the IP addresses configured on this VIF.

- To update an entry, select a configured IP address, change the needed data and click on **UPDATE**.
 - To delete an entry, select a configured IP address and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
9. Click on **OK** to complete the operation. If you configured LAGG, the protocol you chose is displayed in the column **Configuration**.
 10. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them.

Make sure that at least one interface is available, otherwise, you would lose your current connection to SOLIDserver.

Setting up a VIF

A VIF (Virtual Interface) allows to set a number of configurations in a virtual container. Via this container you can apply or edit a network configuration including embedded services. Keep in mind that while the procedures below detail how to create, edit or delete a VIF from the page Network configuration, during each procedure you need to make sure that you have at least one operating interface connected to SOLIDserver or you might lose your point of access, and therefore be unable to manage the appliance.

To add a VIF

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. In the menu, click on **+ Add**.
4. Configure the Virtual interface according to your needs. For more details, refer to the sections [Configuring Basic IP Addressing on an Interface](#), [Setting up a VLAN Interface](#), [Setting up an Ethernet Port Failover](#) and [Configuring a VIP](#).
5. Click on **OK** to complete the operation.
6. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

To edit a VIF

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of the interface you want to edit (all interfaces are preceded by an orange dot). The wizard **Configure network virtual interface** opens.
4. Edit the Virtual interface according to your needs. For more details, refer to the sections [Configuring Basic IP Addressing on an Interface](#), [Setting up a VLAN Interface](#), [Setting up an Ethernet Port Failover](#) and [Configuring a VIP](#).
5. Click on **OK** to complete the operation.
6. Right now your configuration is pending. In the menu, select  **Tools** > **Apply configuration** to save your changes or  **Tools** > **Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

To delete a VIF

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of the interface you want to delete (all interfaces are preceded by an orange dot). The wizard **Configure network virtual interface** opens.
4. In the field **Physical interfaces**, select the interfaces to be deleted one by one and click on . The physical interfaces, are now listed in the field **Available physical interfaces**.
5. Click on **NEXT**. The **IPv4 address configuration** is displayed.
6. In the **IP addresses list**, select the configured IP address(es) one by one. The configuration fields appear.
7. Click on **DELETE**. The IP address is no longer listed in the field.
8. Click on **NEXT**. The **IPv6 address configuration** is displayed. Repeat the steps 5 and 6.
9. Click on **OK** to complete the operation.
10. Right now your configuration is pending. In the menu, select  **Tools** > **Apply configuration** to save your changes or  **Tools** > **Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

Configuring a Media Interface

By default, SOLIDserver automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between the 10/100Base-T and 10/100/1000Base-T ports and the Ethernet ports on a connecting switch. It is usually unnecessary to change the default auto-negotiation setting; however, you can manually configure connection settings for a port if necessary.

Set the media interface

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Network configuration**. The page **Network configuration** opens.
3. Click on the name of **Physical interface** of your choice (it is attached to a VIF or located under the Unused interfaces). The wizard **Network interface configuration** opens.
4. In the drop-down list **Media**, select the speed and duplex to be applied to the physical interface you clicked on. By default, the *autoselect* option is selected, it is automatically selected by SOLIDserver according to your network configuration.
5. Click on to complete the operation.
6. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Chapter 7. Configuring the Services

This chapter details most of the services embedded in SOLIDserver and available on the page **Services configuration**:

- [Handling Services](#) allows to enable/disable and start/stop all the services embedded in SOLIDserver.
- [Configuring the SSH Remote Account](#) allows to set up the details of the connection to SOLIDserver via a Secure Shell (SSH) client.
- [Changing the SFTP/SCP/RSYNC User Account Password](#) allows to edit the *xfer* account password used by the protocols SFTP, SCP and RSYNC¹.
- [Managing the TFTP Upload Authorizations](#) allows to deliver Trivial File Transfer Protocol (TFTP) services in order to send boot and configuration files to DHCP/BOOTP clients (such as IP phones, thin clients, bootless stations).
- [Configuring the SMTP Relay](#) allows to configure the host relay that SOLIDserver uses to send emails via Simple Mail Transfer Protocol (SMTP).
- [Configuring DNS Guardian](#) allows to configure the listening interfaces and enable the service DNS Guardian if your license includes it.
- [Configuring GSLB Server](#) allows to configure the listening interfaces and enable the service GSLB server if your license includes it.
- [Configuring the HTTPS Certificate](#) allows to handle SOLIDserver Apache certificates.
- [Downloading the DNS/DHCP/DHCPv6 Configuration File](#): allows to retrieve all DHCP and DNS configuration files.
- [Managing the SNMP Service](#): allows to monitor SOLIDserver performances and load through the SNMP protocol.

Keep in mind that:

- **All the services must be set at the same time to prevent any management problems.** To this end, we strongly recommend configuring Network Time Protocol (**NTP**) servers to update SOLIDserver timer. For more details, refer to the chapter [Configuring the Time and Date](#).
- **You can set the services of remote appliances** via the drop-down list *SOLIDserver*, whether you simply manage an appliance remotely or configured it in high availability. For more details, refer to the chapter [Centralized Management](#).

¹SFTP stands for Secure File Transfer Protocol also known as *SSH File Transfer Protocol*. SCP stands for Secure Copy. RSYNC stands for Remote Synchronization.

Handling Services

SOLIDserver allows you to completely disable a network service. While a network service is disabled, it cannot run. Once a network service is enabled, its state is automatically updated after having applied the configuration. To sum up, a user can easily handle the embedded services: enabling/disabling and starting/stopping every service provided by SOLIDserver.

Enabling or Disabling a Service

Before enabling or disabling a service, keep in mind that this operation impacts the service:

- **Disabling a service automatically stops it.**
- **Enabling a service automatically starts it**

To enable/disable a service

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, look for the service of your choice.
4. In the column **Enabled**:
 - a. To enable the service, click on the link **Disabled**. The wizard opens.
 - b. To disable the service, click on the link **Enabled**. The wizard opens.
5. Click on to complete the operation.
6. Right now your configuration is pending. In the menu, select  **Tools** > **Apply configuration** to save your changes or  **Tools** > **Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

Starting or Stopping a Service

Before starting or stopping a service, keep in mind that:

- **Once a service is disabled, it cannot be started.**
- **A disabled service is automatically stopped, so you can only stop an *Enabled* service.**

To start/stop a service

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, look for the service of your choice.
4. In the column **Running**:

- a. To start the service, click on the link **Stopped**. The wizard opens.
 - b. To stop the service, click on the link **Started**. The wizard opens.
5. Click on to complete the operation.

Configuring the SSH Remote Account

SOLIDserver can be accessed remotely through an SSH connection using the account *admin* if you enabled the SSH remote console access for SOLIDserver administration. For more details regarding how to enable/disable it, refer to the section [Handling Services](#).

By default, the account *admin* is set with the password is set to *admin*. This account cannot be edited but its password and the password level of security can be edited, as detailed in the sections below.

On the appliance, an SSH shell session is available on SOLIDserver file system. If you update configuration files directly, you can disturb or prevent SOLIDserver from running. Only administrators should use this configuration mode, by default *admin* is the only account that can access SOLIDserver via SSH.

Keep in mind that you can configure SOLIDserver to allow LDAP/RADIUS authentication for SSH connections. For more details, refer to the appendix [Using Remote Authentication for SSH Connections to SOLIDserver](#).

Changing the SSH Remote Access Password

By default the *admin* account password is set to *admin*.

To change the SSH password

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on the link **Account: admin**. The wizard opens.
4. Fill in the password of your choice, in accordance with the level of security you chose, in the fields **New password** and **Confirm password**.
5. Click on to complete the operation.

Changing the SSH Password Level

You can set up the password security level of your choice on the services that use a shell connection through a registry key. There are 3 levels of security:

1. **Low**: the password can contain any character and as few as you want. There is no password restriction, you could set a password with one character.
2. **Medium**: the password requires at least 8 characters, it can be any character.
3. **High**: the password requires at least 8 characters, among which at least 2 must be digits and at least 2 must be special characters (for example: `!, #, @...`).

To change the SSH password security level

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in *module.system.ssh_password*. Only this key is listed.
4. In the column **Value**, click on the digit link. The wizard **Registry database Edit a value** opens. By default, the password level is set to *1*.
5. Set the value of your choice, either *1* (low), *2* (medium) or *3* (high). For more details regarding the levels, refer to the details above this procedure.
6. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Changing the SFTP/SCP/RSYNC User Account Password

The *xfer* account manages the SFTP, SCP and RSYNC services. By default there is no password applied to *xfer* account, so you need to set a password and activate the account to be able to access these services through a shell connection.

To set the xfer account password

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on the link **Account: xfer**. The wizard opens.
4. In the field **New password**, type in the password of your choice, in accordance with the level of security you chose.
5. In the field **Confirm password**, type in the password again.
6. Click on **OK** to complete the operation.

The *xfer* account is not enabled and disabled like the services. Only one wizard allows to enable and disable the account that manages the SFTP, SCP and RSYNC protocols.

To enable/disable the xfer account

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Enabled** of the Account: xfer line, click on **Disabled** or **Enabled**. The wizard opens.
4. Click on **OK** to complete the operation. The account is now marked as **Enabled** or **Disabled**.

Managing the TFTP Upload Authorizations

You can download and upload files through the Trivial File Transfer Protocol (TFTP). From the GUI, you can enable or disable the service. For more details, refer to the section [Handling Services](#) above.

You can enable uploads from remote appliances to SOLIDserver GUI. The uploaded files and files available for download are listed on a dedicated page of the page *Local files listing*. For more details, refer to the section [Managing Files from the Local Files Listing](#).

To enable TFTP uploads to SOLIDserver

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, under TFTP server, click on the link **Upload Authorization: Disabled**. The wizard **TFTP File Upload Authorization** opens.
4. Click on to complete the operation. The report opens and closes. The TFTP **Upload Authorizations** status is now **Enabled**.

Once the uploads are enabled, following the procedure above disables them.

Configuring the SMTP Relay

SOLIDserver provides SMTP (Simple Mail Transfer Protocol) to allow you to add/edit the host relay on emails directly sent through the appliance.

To configure an outgoing mail server

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on **SMTP relay**. The wizard **SMTP Relay Configuration** opens.
4. In the Outgoing mail server, fill in the valid FQDN or the IPv4 address of the server.
5. Click on to complete the operation.
6. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on to complete the operation.

You can also change the source email address of the outgoing mails and alerts notifications. Note that you can only edit the source mail addresses locally.

To change the default source mail

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, under the **Mail - SMTP** line, click on **Default source mail: <mail-address>**. The wizard **Source mail configuration** opens.
4. In the field **Default mail**, type in the email address of your choice.
5. Click on to complete the operation. The new address has now replaced the default address in the list.

To change the alert source mail

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, under the **Mail - SMTP** line, click on **Alert source mail: <mail-address>**. The wizard **Source mail configuration** opens.
4. In the field **Alert mail**, type in the email address of your choice.
5. Click on to complete the operation. The new address has now replaced the default address in the list.

Configuring the HTTPS Certificate

During the first boot, SOLIDserver generates a self-signed certificate used by default to connect to SOLIDserver. As it is not signed by a Certificate Authority (CA), it is not trusted by your web browser and warning messages appear.

You can add or import a certificate and use it as *SSL certificate* as detailed in the procedure below. For more details on certificates import and addition, refer to the section [Managing the HTTPS Certificate](#) in the chapter *Maintenance*.

Note that **the SSL certificate is unique to each SOLIDserver appliance**. So if you want to use a certificate and you are managing remote appliances or appliances in High Availability: use the drop-down list *SOLIDserver* to make sure you are setting each appliance with its own SSL certificate.

To choose an HTTPS certificate

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, select the appliance of your choice.
4. Under the line *HTTP webserver*, click on the link **SSL Certificate**. The wizard **Change the current SSL certificate** opens.
5. In the drop-down list **SSL Certificate**, select the certificate of your choice. By default, the certificate *Apache SSL Cert Base* is available and selected.

6. Click on **OK** to complete the operation.
7. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

Configuring DNS Guardian

If your license includes DNS Guardian, you can configure the listening network interface(s) and enable the dedicated service.

Keep in mind that to enable and configure the service DNS Guardian your appliance must have at least 8 GB of RAM. For more details regarding DNS Guardian, refer to the part [Guardian](#).

Note that if your license includes both DNS Guardian and DNS GSLB, you must configure the line *DNS Guardian / GSLB server* as both features rely on the same service.

To configure the listening interfaces of DNS Guardian

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on **DNS Guardian** or **DNS Guardian / GSLB server**. The wizard **DNS Guardian configuration** or **DNS Guardian & GSLB server configuration** opens.
4. In the list **Available interfaces**, select the interface of your choice and click on **+**. The interface is moved to the list **Selected interfaces**.

Each interface is listed *<interface-name> (<MAC-address>)*, whether it is active or not. Only *Intel* network interfaces are listed as no other interface card can be configured for the service.

5. Repeat this action for as many interfaces as you need.

To remove an interface from the list **Selected interfaces**, select it and click on **-**. The interface is moved back to the list **Available interfaces**.

6. Click on **OK** to complete the operation. The report opens and closes.
7. In the column **Name**, look for *DNS Guardian* or *DNS Guardian / GSLB server*.
8. In the column **Enabled**, click on the link **Disabled** to enable the service. The wizard opens.
9. Right now your configuration is pending. In the menu, select **Tools > Apply configuration** to save your changes or **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **OK** to complete the operation.

The DNS must be running as well. Make sure the services *DNS Guardian* and *DNS server (named)* or *DNS Guardian* and *DNS server (unbound)* are both enabled and started.

Configuring GSLB Server

If your license includes DNS GSLB, you can configure the listening network interface(s) and enable the dedicated service.

Keep in mind that to enable and configure the service GSLB server your appliance must have at least 8 GB of RAM. For more details regarding the configuration of applications with a GSLB server, refer to the part [Application](#).

Note that if your license includes both DNS Guardian and DNS GSLB, you must configure the line *DNS Guardian / GSLB server* as both features rely on the same service.

To configure the listening interfaces of GSLB server

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on **GSLB server** or **DNS Guardian / GSLB server**. The wizard **GSLB server configuration** or **DNS Guardian & GSLB server configuration** opens.
4. In the list **Available interfaces**, select the interface of your choice and click on . The interface is moved to the list **Selected interfaces**.

Each interface is listed *<interface-name> (<MAC-address>)*, whether it is active or not. Only *Intel* network interfaces are listed as no other interface card can be configured for the service.

5. Repeat this action for as many interfaces as you need.

To remove an interface from the list **Selected interfaces**, select it and click on . The interface is moved back to the list **Available interfaces**.

6. Click on  to complete the operation. The report opens and closes.
7. In the column **Name**, look for *GSLB server* or *DNS Guardian / GSLB server*.
8. In the column **Enabled**, click on the link **Disabled** to enable the service. The wizard opens.
9. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes or  **Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on  to complete the operation.

The DNS must be running as well. Make sure the services *GSLB server* and *DNS server (named / nsd / unbound)* are both enabled and started.

Managing the SNMP Service

You can manage the SNMP service directly from the user interface. From the module Administration you can:

- Configure the TCP/UDP ports the server listens on.
- Configure the server v1, v2c and v3 profiles allowed to access SOLIDserver SNMP agent. By default, a v1/v2c profile exists with the community string *public*. You can delete it and create custom profiles to secure your system.
- Configure SNMP Traps on the server for a network management platform.

The columns *Running* and *Enabled* on the page Services configuration indicate the server state on the appliance.

Note that SNMPv3 requires a properly configured NTP server. For more details, refer to the section [Configuring NTP Servers](#).

To configure the SNMP server

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. Click on the link **SNMP server**. The wizard **SNMP Server Configuration** opens.
4. In the fields **UDP port** and **TCP port**, type in the port number to communicate with the protocol of your choice. At least one field has to be filled in. By default, the UDP port number used is *161*, you can also use that port with TCP.
5. Click on **NEXT**. The next page opens.
 - a. Configure an SNMP profile following the table below.

Table 7.1. SNMP profiles parameters

Parameters	Description
SNMP version	You can choose either <i>v1/v2c</i> or <i>v3</i> . By default, <i>v1/v2c</i> is selected.
Access	No matter the SNMP version, this field cannot be edited and is by default in <i>Read-only</i> .
<i>v1/v2c</i>	SNMPv1 and SNMPv2c are simple request/response protocols. SNMPv2c support includes a bulk-retrieval mechanism and more detailed error message reporting to management stations. If you select <i>v1/v2c</i> , the following fields are displayed: <i>Access</i> , <i>Community</i> and <i>SNMP restriction</i> .
<i>v3</i>	SNMPv3 uses the security features providing secure access to devices. If you select <i>v3</i> , the following fields are displayed: <i>Access</i> ; <i>Users</i> , <i>Key</i> , <i>Protocol</i> (authentication fields) and <i>Key</i> and <i>Protocol</i> (privacy fields).
Community	Type in the community string that would act as a password to access the SNMP agent.
SNMP restriction	Type in the source of the SNMP. It can be one IP address, several IP addresses separated by a space or a default value.
User	Type in the login used for authentication.
Level	You can choose either a <i>noauth</i> , <i>auth</i> or <i>priv</i> security level.
Key	Type in the authentication passphrase (i.e. password), it must contain at least 8 characters.
Algorithm	You can select either the <i>MD5</i> ^a or the <i>SHA</i> ^b algorithm. By default, <i>MD5</i> is selected.
Key	Type in a privacy passphrase. If the privacy passphrase is not specified, it is assumed to be the same as the authentication passphrase. This field is optional
Protocol	You can select either the <i>DES</i> ^c or <i>AES</i> ^d algorithm. By default, <i>DES</i> is selected.

^aMD5 Message-Digest algorithm.

^bSecure Hash Algorithm.

^cData Encryption Standard.

^dAdvanced Encryption Standard.

- b. When the access configuration is complete, click on **ADD**. The profile is moved to the **SNMP access list**.

- To update an entry, select the SNMP profile of your choice, change the data according to your needs and click on **[UPDATE]**.
 - To delete an entry, select the SNMP profile of your choice and click on **[DELETE]**.
 - To discard the latest changes, click on **[CANCEL]**.
- c. Repeat these operations for as many SNMP profiles as you need.
6. Click on **[NEXT]**. The last page of the wizard opens.
- a. You can set an SNMP trap configuration, following the table below.

Table 7.2. SNMP trap configuration

Parameters	Description
Send Trap v1	You can choose to enable an agent ^a to send a trap notifying the management station of significant events through the SNMP v1 protocol. By default, Yes is selected.
Send Trap v2	You can choose to enable an agent to send a trap notifying the management station of significant events through the SNMP v2 protocol. By default, Yes is selected.
Send Trap Inform	You can choose to enable routers to send inform requests to SNMP managers. By default, Yes is selected.
Host	Type in the IP address of the computer that listens to the network and catches the trap.
Port	You can define through which port the host that should catch the trap. This field is optional.
Community	Type in the community string that would act as a password to access the SNMP agent.

^aDetails regarding agent can be found in the Management Information Base (MIB)

- b. When your configuration is complete, click on **[ADD]**. The profile is moved to the **Trap list**.
- To update an entry, select the SNMP trap of your choice, change the data according to your needs and click on **[UPDATE]**.
 - To delete an entry, select the SNMP trap of your choice and click on **[DELETE]**.
 - To discard the latest changes, click on **[CANCEL]**.
- c. Repeat these operations for as many traps as you need.
7. Click on **[OK]** to complete the operation.
8. Right now your configuration is pending. In the menu, select **✖. Tools > Apply configuration** to save your changes or **✖. Tools > Rollback configuration** to discard them. The corresponding wizard opens, click on **[OK]** to complete the operation.

Downloading the DNS/DHCP/DHCPv6 Configuration File

The Services configuration page allows you to download the current DNS (named.conf), DHCP (dhcpd.conf), DHCPv6 (dhcpd6.conf), NSD (nsd.conf) or Unbound (unbound.conf) configuration file of the appliance of your choice: whether the local one or the configuration of one of the appliances you are managing remotely.

To download the DNS/DHCP/DHCPv6/NSD/Unbound configuration file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the drop-down list **SOLIDserver**, under the menu, select the appliance for which you want to download the configuration file.
4. In the menu, select **Tools > Download configuration file**. The wizard **Download configuration file** opens.
5. In the drop-down list **Configuration file**, select *DNS*, *DHCP*, *DHCP V6*, *NSD* or *Unbound* according to your needs.
6. Click on **OK** to complete the operation. The report opens, the configuration file can be downloaded from the page **Local files listing** available from the page *Admin Home*.
7. Click on **DOWNLOAD** to download the file before closing the wizard.
8. Click on **CLOSE**. The wizard closes and the page *Services configuration* is visible again. Note that the report is generated and stored on the page *Local Files Listing*. For more details regarding reports, refer to the chapter [Managing Reports](#).

Chapter 8. Managing the Licenses

Once you added a license, you can renew it for one or several appliances or delete it locally at any point.

The page **Centralized Management** provides information regarding each appliance license and maintenance period in dedicated columns. For more details, refer to the table [The default columns on the page Centralized Management](#).

Renewing a License

Renewing a license is necessary if you want to manage more services or if you are notified in the GUI in a banner above the top bar or in the gadget *System Information* that a temporary license or your maintenance period is expiring.

Before renewing a license, note that:

- For a local appliance you must:
 1. Retrieve the request key and send it to EfficientIP via the dedicated portal, as detailed in the section [Requesting a License Key for the Local Appliance](#).
 2. When you receive the new license key, you must add it to your appliance to activate the license as detailed in the section [Activating a License](#).
- For remote appliances you can:
 1. Retrieve locally the request key on each remote appliance, connect to the management appliance to export them all at once, and send them to EfficientIP on the dedicated portal, as detailed in the section [Requesting a License Key for Remote Appliances](#).
 2. When you receive the new license keys, you must go to the page *All Centralized Management* and import the license keys to activate the license on all the remote appliances at once, as detailed in the section [Activating a License](#).

Requesting a License Key for the Local Appliance

From the page *License* of any appliance, you can retrieve the license request key and send it to EfficientIP.

To request a new license key

Only users of the group *admin* can perform this operation.

1. Retrieve the request license key.
 - a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **System**, click on **License**. The page **License** opens.
 - c. In the menu, select **+ Add > Request license**. The wizard **Request license** opens.
 - d. Copy the content of the field **Request key**, it is required when you fill out the request license form.
 - e. Click on to close the wizard.

2. Send the request key to Efficient IP.
 - a. Go to the page <http://www.efficientip.com/license-request/> and fill out the *Request Your License* form.
 - b. In the fields **FIRST NAME**, **LAST NAME**, **EMAIL**, **COMPANY**, **PHONE NUMBER** and **COUNTRY NAME**, specify your contact details. All these fields are required.
 - c. In the field **LICENSE PERIOD REQUEST**, specify the desired license length: *1 month*, *2 months*, *3 months*, *6 months* or *Permanent*. This field is required.
 - d. If you selected *Permanent*, you must fill in the field **CONTRACT NUMBER (IF PERMANENT LICENSE)**.
 - e. In the field **REQUEST KEY**, paste your request key or the content of your request key file. This field is required.
 - f. In the field **NUMBER OF EXTERNAL MANAGED SERVERS (MVSM, IF ANY)**, specify the total number of servers - DNS/DHCP/... - you intend to manage from SOLIDserver.
 - g. In the section **OPTIONAL MODULE**, tick all the optional modules you might need: *NETCHANGE*¹, *DEVICE MANAGER*, *SPX* or *DNS GUARDIAN*.
 - h. If relevant, fill in the field **IF REQUESTER IS OTHER THAN THE END CUSTOMER, PLEASE PROVIDE YOUR CONTACT INFORMATION (NAME, COMPANY, EMAIL, PHONE)**.
 - i. Click on to send us your information.

Once EfficientIP has answered your request and sent you a license key, you can renew your licence as detailed in the section [Activating a License](#).

Requesting a License Key for Remote Appliances

From the page *Centralized Management* of a management appliance, you can export the license request key of all remote appliances and send them at once to EfficientIP.

To request the license request key of all the remote appliances

Only users of the group *admin* can perform this operation.

1. Retrieve the request license key.
 - a. Connect to the GUI of the remote appliance of your choice.
 - b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
 - c. In the section **System**, click on **License**. The page **License** opens.
 - d. In the menu, select **+ Add > Request license**. The wizard **Request license** opens.
 - e. Copy the content of the field **Request key**, it is required when you fill out the request license form.
 - f. Click on to close the wizard.
 - g. Repeat the steps a-f for all the relevant remote appliances.

¹If you do not tick this box, you are using NetChange basic options, or NetChange-IPL.

2. Export all the license keys from the management appliance.
 - a. Connect to the management appliance GUI.
 - b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
 - c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
 - d. In the menu, select **Tools > Export license requests**. The wizard opens.
 - e. Read the **License Agreement** and click on **NEXT**. The page **Export license requests** opens.
 - f. Click on **OK** to complete the operation. The page **Report** opens.
 - g. Click on **DOWNLOAD** to save the file locally. When the file download is complete, the page *Centralized Management* is visible again.

3. Send the request key to Efficient IP.
 - a. Go to the page <http://www.efficientip.com/license-request/> and fill out the *Request Your License* form.
 - b. In the fields **FIRST NAME**, **LAST NAME**, **EMAIL**, **COMPANY**, **PHONE NUMBER** and **COUNTRY NAME**, specify your contact details. All these fields are required.
 - c. In the field **LICENSE PERIOD REQUEST**, specify the desired license length: *1 month*, *2 months*, *3 months*, *6 months* or *Permanent*. This field is required.
 - d. If you selected *Permanent*, you must fill in the field **CONTRACT NUMBER (IF PERMANENT LICENSE)**.
 - e. In the field **REQUEST KEY**, paste your request key or the content of your request key file. This field is required.
 - f. In the field **NUMBER OF EXTERNAL MANAGED SERVERS (MVSM, IF ANY)**, specify the total number of servers - DNS/DHCP/... - you intend to manage from SOLIDserver.
 - g. In the section **OPTIONAL MODULE**, tick all the optional modules you might need: *NETCHANGE²*, *DEVICE MANAGER*, *SPX* or *DNS GUARDIAN*.
 - h. If relevant, fill in the field **IF REQUESTER IS OTHER THAN THE END CUSTOMER, PLEASE PROVIDE YOUR CONTACT INFORMATION (NAME, COMPANY, EMAIL, PHONE)**.
 - i. Click on **SUBMIT** to send us your information.

Once EfficientIP has answered your request and sent you license keys, you can renew your licenses as detailed in the section [Activating a License](#).

Activating a License

Once you received the license key(s), you must activate it:

- For local appliances, you can add the license key on the page *License* or import it on the page *Centralized Management*.

²If you do not tick this box, you are using NetChange basic options, or NetChange-IPL.

- For remote appliances, you import all the license keys at once on the page *Centralized Management*.

To activate a license locally

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **License**. The page **License** opens.
3. In the menu, select **+ Add > License**. The wizard opens.
4. Read the **License Agreement** and click on **NEXT**. The page **Import licenses** opens.
5. In the field **License(s)**, paste the license key.
6. Click on **OK** to complete the operation.

To activate one or more licenses at once

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
3. Tick the appliance(s) of your choice.
4. In the menu, select **Tools > Import licenses**. The wizard opens.
5. Read the **License Agreement** and click on **NEXT**. The page **Importing licenses** opens.
6. In the drop-down list **Import type**, choose the import method.
 - a. To paste the key(s) yourself, select **Manual copy** and, in the field **License(s)**, paste the license key(s). If you paste several keys, enter two line breaks between each key.
 - b. To look for the key(s) on your computer, select **File** and click on **BROWSE** to select the *.txt* file containing the license key(s).

Note that the license you import automatically overwrites the current license on the relevant appliance(s).

7. Click on **OK** to complete the operation.

Deleting a License

At any point, you can delete a license. Keep in mind that:

- Deleting a license must be done locally.
- Deleting the license also deletes the maintenance.

Note that renewing a license does not require deleting the current license. When you activate a license locally, it automatically replaces the license currently installed.

To delete a license

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **License**. The page **License** opens.
3. In the panel **Current license**, click on **DELETE**. The wizard **Delete the license** opens.
4. Click on **OK** to complete the operation. The panel **Current license** is now empty. In the panel **System information**, the *License type* indicates *No license installed*.

Part III. Imports and Exports

SOLIDserver supports many imports and exports methods in almost all the modules..

- [Importing Data from a CSV File](#) details how to import or reimport data from a CSV file in the modules IPAM, DHCP, DNS, NetChange, Devic Manager, VLAN Manager and Administration.
 - [Importing IPAM Data](#) details how to import VitalQIP and Nortel NetID data to the module IPAM.
 - [Importing DHCP Data](#) details how to import ISC, Alcatel-Lucent VitalQIP, Microsoft, Infoblox, MetaIP and Nortel NetID configuration files to the module DHCP.
 - [Importing DNS Data](#) details how to import BIND and VitalQIP zones from an archive file to the module DNS.
 - [Exporting Data](#) details how to export data from any module to CSV, HTML, XML, Excel or PDF files.
-

Chapter 9. Importing Data from a CSV File

You can massively import data from CSV formatted files in the modules IPAM, DHCP, DNS, NetChange, Device Manager, VLAN Manager, VRF and Administration. To import specific IPAM, DNS or DHCP configurations in a different format, refer to the import chapter in each module.

Before importing CSV files, keep in mind that:

- **The user importing the data must have the appropriate administrative rights**

For instance, importing addresses into a terminal network implies that the user has administrative rights over said network.

- **The resources you import must have a unique name**

You can import as many resources as you need as long as they all have a unique name.

- **An import is generated one page at a time**

If you are importing zones from the page *All zones* in the DNS, you only import the zones themselves but not the RRs they contain.

- **The object parameters that you can import correspond to the columns of the page**

When you import objects, you can import all or some of their properties. That way, for instance, you can import a zone and specify its view, this allows you to set up your DNS hierarchy more easily.

- **An import can overwrite the existing page data**

The last step in the import wizard allows to overwrite, or not, all the existing parameters of the resource except for its name.

- **If the page does not have the menu Import you cannot import data**

You can import data from almost any page. In the menu, CSV <data> allows to import CSV files. The list of pages where you can import data is available in each module-dedicated import section below.

- **Your classes may edit the fields available in the wizard**

The import procedure only describes the fields that are displayed by default for each resource. During each import, the number of fields that might appear or be required depends on the classes you or your administrator might have configured for each resource.

- **Any object parameter value is imported with the Inheritance property forced to *Set* or *Inherit***

This allows to respect the configuration of the Inheritance property of an object parameter:

- It is forced to *Set* if the parameter is not configured on the parent object or if it has a different value.
- It forced to *Inherit* if it is configured on the parent object and has the same value.

The Import Wizard

During an import, the wizard displays a set of pages that you should configure properly. Once you selected the CSV file you want to import, the page **CSV fields association**, opens.

Figure 9.1. The first page of the import wizard

- ❶ This section allows to specify the CSV import file details. They can be configured and saved as templates to speed up the checking process. Its fields are detailed in the table [CSV file basic parameters](#).
- ❷ This section contains some parameters (columns) that you can include in your import.

The the first section of the import wizard is common to all objects and can be configured as follows.

Table 9.1. CSV file basic parameters

Parameter	Description
Delimiter	Select the data delimiter of your choice (a comma, a semi-colon or a tab) in the drop-down list. The comma is selected by default.
Enclosure	Select the data enclosure of the text (a single quotation or a double quotation mark) in the drop-down list. The double quotation mark is selected by default.
Input format	Select the input format of your data (<i>UTF-16</i> , <i>ASCII</i> or <i>UTF-8^B</i>) in the drop-down list. The <i>ASCII</i> format is selected by default.
Skip first line	Select <i>Yes</i> or <i>No</i> in the drop-down list depending on your needs. Skipping the first line avoids importing the columns title. <i>Yes</i> is selected by default.
Template	Select <i>None</i> if you do not want to save your parameters in a template. Select <i>New template</i> to save your parameters as a CSV file export template. This drop-down list also contains the existing templates that you can reuse.
Template name	If you selected <i>New template</i> , name it in this field.

Parameter	Description
Save template	This box is visible if you did not create a template or if you selected an existing template. Tick it if you want to save the changes made to an existing template.

^aThe UTF-8 input format is necessary to successfully import CSV files containing accents.

On the next page, **Class parameters**, there are many drop-down lists as there are existing class parameters for the resource you are importing. None of the lists are required, they allow to make a specific and detailed import if your parameters match a class in the database or if they can be interpreted. Any class parameter that does not correspond to a class in the database is not displayed in the GUI once imported.

Finally, on the last page, **CSV import parameters**, a few options are available.

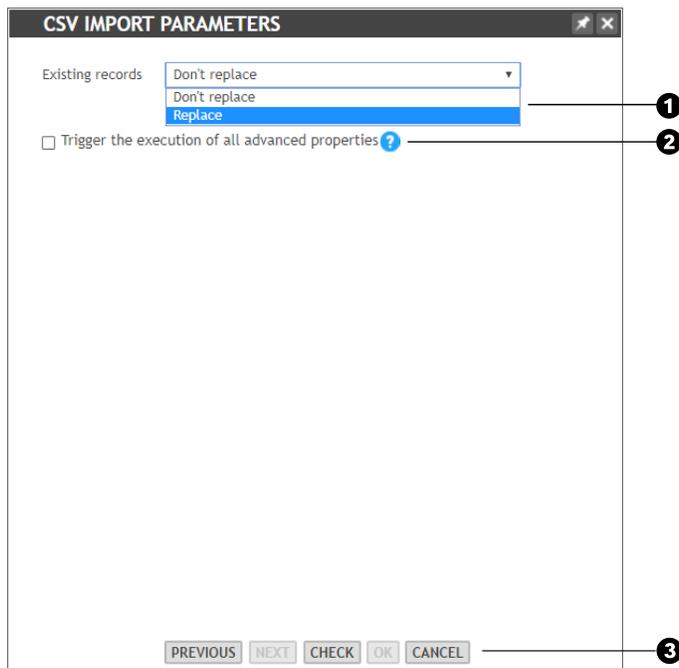


Figure 9.2. The consistency check page of the import wizard

- ❶ This drop-down list **Existing records** allows to replace the existing entries with the data of the CSV file you are importing.
- ❷ The box **Trigger the execution of all advanced properties** allows to force the DNS, DHCP and VLAN advanced property mechanisms during the import of networks, pools and addresses. For more details, refer to the section [Importing Data to the IPAM](#)
- ❸ The button **CHECK** performs a data validity check of the content of the CSV file. The last pages of the wizard provide a **Report**: a data validity report and an import report.

Importing Data to the IPAM

Within the IPAM module, you can import data on every page, except the pages *All deleted IP addresses* (neither in IPv4 nor in IPv6) and *All policies*. The table below details where you can import them within the module.

Table 9.2. IPAM pages where you can import CSV files

IPAM page	Objects that can be imported	Option name in the menu  Import
All spaces	Spaces	CSV spaces
	Block-type networks	CSV networks (block)
	Subnet-type networks	CSV networks (subnet)
	Pools	CSV pools
	Addresses	CSV addresses
	IPv6 block-type networks	CSV networks (block v6)
	IPv6 subnet-type networks	CSV networks (subnet v6)
	IPv6 pools	CSV pools (v6)
	IPv6 addresses	CSV addresses (v6)
	SPX allocated networks	SPX allocated networks
	SPX assigned networks	SPX assigned networks
	IPv6 SPX allocated networks	SPX allocated networks (v6)
	IPv6 SPX assigned networks	SPX assigned networks (v6)
All networks	Block-type networks	CSV networks (blocks)
	Subnet-type networks	CSV networks (subnets)
	Pools	CSV pools
	IP Addresses	CSV addresses
	SPX allocated networks	SPX allocated networks
	SPX assigned networks	SPX assigned networks
All networks (v6)	IPv6 block-type networks	CSV networks (block v6)
	IPv6 subnet-type networks	CSV networks (subnet v6)
	IPv6 pools	CSV pools (v6)
	IPv6 addresses	CSV addresses (v6)
	SPX IPv6 allocated networks	SPX allocated networks (v6)
	SPX assigned networks (v6)	SPX assigned networks (v6)
All pools	Pools	CSV pools
	Addresses	CSV addresses
All pools (v6)	IPv6 pools	CSV pools (v6)
	IPv6 addresses	CSV addresses (v6)
All addresses	Addresses	CSV addresses
All addresses (v6)	IPv6 addresses	CSV addresses (v6)

Before importing, keep in mind that:

- The advanced properties can automate the creation of entries in your database after the import. If all the advanced properties are activated, importing IPAM data may automatically update the DHCP and DNS databases. If you do not want your import to impact other modules, edit the *Internal module setup* before importing IPAM data. For more details, refer to the chapter [Managing Advanced Properties](#).
- Other advanced options, like the IPv4 to IPv6 transition, the IPAM / Device Manager interaction properties and the IPAM/VLAN interaction properties, may also update your database automatically. For more details, refer to the chapters [Setting Up a Transition From IPv4 to IPv6](#), [Managing the Interaction with the IPAM](#) and [Managing the IPAM/VLAN Interaction](#).

To import VitalQIP or Nortel NetID data, refer to the chapter [Importing IPAM Data](#).

To import SPX allocated and assigned networks, whether RIPE or APNIC, refer to the sections [Importing SPX Allocated Networks](#) and [Importing SPX Assigned Networks](#).

Importing Spaces

When importing space(s), only the field **Name** is required. The other parameters are optional and can be left empty.

To import spaces through a CSV file

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. In the menu, select **Import > CSV spaces**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

Only the field **Name** is required. All fields are detailed in the table below.

Table 9.3. Space import parameters

Parameter	Description
Name	The space name. This field is required.
Description	The space description. This field is optional.
Parent space	The space parent space (VLSM), if relevant. This field is optional.
Class parameters	The space-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Class name	The space class name. This field is optional.

7. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **NEXT**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If

you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.

13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All spaces**. The spaces are now listed.

Importing Networks

You can import block-type networks or subnet-type networks in IPv4 or IPv6.

Importing IPv4 Block-type Networks

Block-type networks must be imported into a space.

When importing a file containing any type of network, the field **First address** and one of the fields specifying the network size are required: **Last address**, **Netmask**, **Prefix** or **Size** must be specified.

Note that the procedure below is based on an import made on the page **All networks**. You can also import networks in a specific space, in which case the field **Space name** is not displayed in the wizard.

To import IPv4 block-type networks through a CSV file

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV networks (blocks)**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The field **First address** and any field that indicates the network size (**Last address**, **Netmask**, **Prefix** or **Size**) are required. All fields are detailed in the table below.

Table 9.4. Block-type network import parameters

Parameter	Description
First address	The network first address. This field is required.
Last address	
Netmask	The drop-down lists that specify the size of the network(s), the number of IP addresses being imported. You must fill at least one of them.
Prefix	
Size	
Name	The network name. This field is optional.
Class parameters	The network-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.

Parameter	Description
Description	The network description. This field is optional.
Class name	The network class name. This field is optional.
Space name	The name of the space where you want import the network(s). It can be a space listed in the file or an existing space in your database. This field is required.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
13. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
15. Click on **CLOSE** to go back to the page **All networks**. The networks are now listed.

Importing IPv6 Block-type Networks

Block-type networks must be imported into a space.

When importing a file containing any type of network, the field **First address** and one of the fields specifying the network size are required: **Last address**, **Netmask**, **Prefix** or **Size** must be specified.

Note that the procedure below is based on an import made on the page **All networks**. This allows you to import a network from any space into the space that suits your needs. You can also import networks in a specific space, in which case the field *Space name* is not displayed in the wizard.

To import IPv6 block-type networks through a CSV file

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV networks (block v6)**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.

6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **First address**, **Prefix** and **Space name** are required. All fields are detailed in the table below.

Table 9.5. IPv6 block-type network import parameters

Parameter	Description
First address	The first address of the IPv6 network(s). This field is required.
Prefix	The network Prefix. This field is required.
Name	The network name. This field is optional.
Class name	The network class name. This field is optional.
Class parameters	The network-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Description	The network description. This field is optional.
Space name	The name of the space where you want import the network(s). It can be a space listed in the file or an existing space in your database. This field is required.

8. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on [NEXT](#). The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on [CHECK](#). The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the validity report in the specified file format.
13. Click on [OK](#) to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the import report in the specified file format.
15. Click on [CLOSE](#) to go back to the list **All networks**. The networks are now listed.

Subnet-type Networks Import Specificities

Several options and fields are only available for the import of subnet-type networks.

Use best space

This option is available in the drop-down list *Space name*, if you import IPv4 subnet-type networks on the page *All networks* (rather than within a specific space or block-type network).

It allows to import the networks into the space containing the smallest network able to receive them.

Note that importing subnet-type networks into a space containing no matching block-type network places them in a container *Orphan Networks*. Later on, you can create a block-type network large enough to contain them. For more details, refer to the chapter [Adding Networks](#).

VLSM space name

This option allows to import subnet-type networks in a specific level of a space-based VLSM organization. Note that to import subnet-type networks in a VLSM configuration:

- To import subnet-type networks in a space and replicate them as block-type networks in one of its sub-spaces, in the field *Space name*, select the parent space and in the field *VLSM space name*, select the sub-space.
- You cannot select the option *Use best space* in drop-down list *Space name*.
- You cannot specify a *VLSM space name* and tick the box *Imbricated networks* in one import. If you do, the *VLSM space name* prevails and the option *Imbricated networks* is ignored. You can configure them in two separate imports if you need them both in your network configuration.

Imbricated networks

This box allows to import a network-based VLSM organization. Tick it to import non-terminal subnet-type networks and all the terminal and non-terminal networks they contain.

Note that you cannot tick the box and specify a *VLSM space name* in one import. If you do, the *VLSM space name* prevails and the option *Imbricated networks* is ignored. You can configure them in two separate imports if you need them both in your network configuration.

If you want to import an organization without ticking the box, the subnet-type networks are imported in a container *Orphan Networks* in the order saved in the CSV file. The first are imported, the rest is considered overlap and is not imported.

If you import an organization and tick the box, the receiving container shapes the import behavior:

- If the selected *Space name* does not contain any block-type network to receive it, the first non-terminal subnet-type network becomes a block-type network.
- If the selected *Space name* contains block-type networks:
 - If an existing block-type network is bigger than the first non-terminal subnet-type network, the whole hierarchy is created within the block-type network if there is enough space available. Otherwise, only the subnet-type networks that fit in the block-type network are imported.
 - If an existing block-type network is the same size as the first non-terminal subnet-type network, the first non-terminal subnet-type network is ignored. The subnet-type networks it contains are imported in the block-type network if it can receive them. Otherwise, only the subnet-type networks that fit in the block-type network are imported.

Importing IPv4 Subnet-type Networks

Subnet-type networks can be imported into a block-type network or directly into a space.

When importing a file containing any type of network, the field **First address** and one of the fields specifying the network size are required: **Last address**, **Netmask**, **Prefix** or **Size** must be specified.

If you import IPv4 subnet-type networks on the page *All networks* (rather than within a specific space or block-type network), the option *Use best space* is available in the drop-down list **Space name**. It allows to put the content of the CSV file into the space containing the smallest block-type network able to receive the subnet-type network(s). Other options are detailed in the section [Subnet-type Networks Import Specificities](#).

To import IPv4 subnet-type networks through a CSV file

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV networks (subnet)**. The wizard **Import a CSV file** opens.
4. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **[NEXT]**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **Address**, **Name**, **Space name** and one field specifying the network size (**Netmask**, **Prefix** or **Size**) are required. All fields are detailed in the table below.

Table 9.6. IPv4 subnet-type network import parameters

Parameter	Description
Address	The network start address. This field is required.
Netmask	The drop-down lists that specify the size of the network(s), the number of IP addresses being imported. You must fill at least one of them.
Prefix	
Size	
Name	The network name. This field is required.
Network is terminal	<p>The VLSM status of the network(s), terminal or non-terminal. This field is optional.</p> <p>By default, if you import a network hierarchy, the last imbricated network is considered terminal even if it is not. You can select the column Terminal to ensure that the networks remain terminal, or non-terminal, during the import. This parameter is only taken into account if its value is <i>0</i> or <i>1</i>.</p> <p>Note that if you tick the box Imbricated networks, you must leave this field empty.</p>
Class parameters	The network-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Space name	<p>The name of the space where you want to import the network(s). It can be a space listed in the file, an existing space in your database or you can use the option <i>Use best space</i>. This field is required.</p> <p>Use best space: Select this option to import subnet-type networks in the space managing the smallest networks that can receive them.</p> <p>Note that you cannot select this option if you want to specify a VLSM space name.</p>
VLSM space name	If you set up a space-based VLSM organization, select the sub space that uses the subnet-type network you are importing as a block-type network. This field is optional.

Parameter	Description
	Note that if you tick the box Imbricated networks or select the option <i>Use best space</i> in the drop-down list Space name , you must leave this field empty.
Class name	The network class name. This field is optional.
Imbricated networks	<p>Tick this box if you want to import a hierarchy of non-terminal and terminal (subnet-type) networks.</p> <p>Note that if you tick this box, you should leave blank the fields Network is terminal and VLSM space name.</p>

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **CSV import parameters** opens:
 - a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
 - b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include properties impacting the IPAM, DNS, DHCP and VLAN Manager. For more details, refer to the section [Network Advanced Properties](#) or to the chapter [Managing the IPAM/VLAN Interaction](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the subnet-type networks of your choice and in the menu select **Tools > Expert > Initialize rules**. This operation also triggers the replication on the objects they contain.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All networks**. The networks are now listed.

Importing IPv6 Subnet-type Networks

Subnet-type networks can be imported into a block-type network or directly into a space.

When importing a file containing any type of network, the field **First address** and one of the fields specifying the network size are required: **Last address**, **Netmask**, **Prefix** or **Size** must be specified.

Note that in IPv6, the option *Use best space* is not available in the drop-down list **Space name**. Available options are detailed in the section [Subnet-type Networks Import Specificities](#).

To import IPv6 subnet-type networks through a CSV file

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV networks (subnet v6)**. The wizard **Import a CSV file** opens.
4. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **[NEXT]**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **Address**, **Prefix**, **Name** and **Space name** are required. All fields are detailed in the table below.

Table 9.7. IPv6 subnet-type network import parameters

Parameter	Description
Address	The network start address. This field is required.
Prefix	The size of the network(s), the number of IP addresses being imported. This field is required.
Name	The network name. This field is required.
Network is terminal	The VLSM status of the network(s), terminal or non-terminal. This field is optional. By default, if you import a network hierarchy, the last imbricated network is considered terminal even if it is not. You can select the column Terminal to ensure that the networks remain terminal, or non-terminal, during the import. This parameter is only taken into account if its value is <i>0</i> or <i>1</i> . Note that if you tick the box Imbricated networks , you must leave this field empty.
Class parameters	The network-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
VLSM space name	If you set up a space-based VLSM organization, select the sub space that uses the subnet-type network you are importing as a block-type network. This field is optional. Note that if you tick the box Imbricated networks , you must leave this field empty.
Space name	The name of the space where you want import the network(s). It can be a space listed in the file or an existing space in your database. This field is required.
Class name	The network class name. This field is optional.
Imbricated networks	Tick this box if you want to import a hierarchy of non-terminal and terminal (subnet-type) networks. Note that if you tick this box, you should leave blank the fields Network is terminal and VLSM space name .

8. Click on **[NEXT]**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you

need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.

9. Click on **NEXT**. The page **CSV import parameters** opens:
 - a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
 - b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include properties impacting the IPAM, DNS, DHCP and VLAN Manager. For more details, refer to the section [Network Advanced Properties](#) or to the chapter [Managing the IPAM/VLAN Interaction](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the subnet-type networks of your choice and in the menu select **Tools > Expert > Initialize rules**. This operation also triggers the replication on the objects they contain.

10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All networks**. The networks are now listed.

Importing Pools

Keep in mind that you **cannot import pools in an empty space**: to successfully import pools, you need a terminal network that can receive them.

Importing IPv4 Pools

On the page **CSV fields association**, the fields **First address**, **Last address**, **Name** and **Space name** are required.

If you import IPv4 pools on the page *All networks* (rather than within a specific space or block-type network), the option *Use best space* is available in the drop-down list *Space name*. It allows to put the content of the CSV file into the space containing the smallest network able to receive the pool(s).

To import IPv4 pools through a CSV file

1. In the sidebar, go to **IPAM > Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.

3. In the menu, select **⬅️ Import > CSV pools**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **First address**, **Name**, **Space name** and one of fields specifying the pool size (**Last address** or **Size**) are required. All fields are detailed in the table below.

Table 9.8. IPv4 pool import parameters

Parameter	Description
First address	The pool first address. This field is required.
Last address	The drop-down lists that specify the size of the pool(s), the number of IP addresses being imported. You must fill at least one of them.
Size	
Name	The pool name. This field is required.
Read-only	The pool reservation status. This field is optional.
Class name	The pool class name. This field is optional.
Space name	The name of the space where you want import the pool(s). It can be a space listed in the file, an existing space in your database or you can use the option <i>Use best space</i> . This field is required.
	Use best space: Select this option to import pools in the space managing the smallest networks that can receive them.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **CSV import parameters** opens:
 - a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
 - b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include the IPAM to DHCP replication. For more details, refer to the section [Pool Advanced Properties](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the pools of your choice and in the menu select **⚙️ Tools > Expert > Initialize rules**. This operation also triggers the replication on the IP addresses they contain.

10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.

11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All pools**. The pools are now listed.

Importing IPv6 Pools

When importing IPv6 pools, you must specify a space with a terminal network able to receive the pools. On the page **CSV fields association**, the fields **First address**, **Last address**, **Name** and **Space name** are required.

Note that in IPv6, the option *Use best space* is not available in the drop-down list **Space name**.

To import IPv6 pools through a CSV file

1. In the sidebar, go to **IPAM > Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV pools (v6)**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **First address**, **Last address**, **Name** and **Space name** are required. All fields are detailed in the table below.

Table 9.9. IPv6 pool import parameters

Parameter	Description
First address	The pool first address. This field is required.
Last address	The drop-down lists that specify the size of the pool(s), the number of IP addresses being imported. You must fill at least one of them.
Size	
Name	The pool name. This field is required.
Read-only	The pool reservation status. This field is optional.
Class name	The pool class name. This field is optional.
Space name	The name of the space where you want import the pool(s). It can be a space listed in the file or an existing space in your database. This field is required.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you

need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.

9. Click on **NEXT**. The page **CSV import parameters** opens:
 - a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
 - b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include the IPAM to DHCP replication. For more details, refer to the section [Pool Advanced Properties](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the pools of your choice and in the menu select **Tools > Expert > Initialize rules**. This operation also triggers the replication on the IP addresses they contain.

10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All pools**. The pools are now listed.

Importing IP Addresses

Keep in mind that **you can import addresses in an empty space**, they are saved in a container *Orphan Addresses*.

Importing IPv4 Addresses

When importing a file containing IPv4 address(s), on the page **CSV fields association**, the fields **IP address**, **Name** and **Space name** are required.

To import IPv4 addresses through a CSV file

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV addresses**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.

6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **IP address**, **Name** and **Space name** are required. All fields are detailed in the table below.

Table 9.10. IPv4 address import parameters

Parameter	Description
IP address	The IP address. This field is required.
Name	The IP address name. This field is required.
MAC address	The IP address MAC address. This field is optional.
Alias	The alias associated with the IP address(es). This field is optional.
Class parameters	The IP address-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Space name	The name of the space where you want import the address(es). It can be a space listed in the file, an existing space in your database or the option <i>Use best space</i> . This field is required. Use best space: Select this option to import addresses in the space containing a block-type network and subnet-type network that can receive the IP address(es) ^a .
Class name	The class name of the IP address(es) you are importing.

^aTo import IP addresses in a container *Orphan Addresses*, import them on the page All networks or All pools of a specific space.

8. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on [NEXT](#). The page **CSV import parameters** opens:
 - a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
 - b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include properties impacting the IPAM, DNS, DHCP and Device Manager. For more details, refer to the section [IP Address Advanced Properties](#) or to the chapter [Managing the Interaction with the IPAM](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the IP addresses of your choice and in the menu select **Tools > Expert > Initialize rules**.

10. Click on [CHECK](#). The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the validity report in the specified file format.

12. Click on **[OK]** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **[TEXT]**, **[HTML]**, or **[EXCEL]** to download the import report in the specified file format.
14. Click on **[CLOSE]** to go back to the page **All addresses**. The IP addresses are now listed.

Importing IPv6 Addresses

When importing a file containing IPv6 address(s), on the page **CSV fields association**, the fields **IP address**, **Name** and **Space name** are required.

To import IPv6 addresses through a CSV file

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV addresses (v6)**. The wizard **Import a CSV file** opens.
4. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **[NEXT]**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **IP address**, **Name** and **Space name** are required. All fields are detailed in the table below.

Table 9.11. IPv6 address import parameters

Parameter	Description
IP address	The IP address. This field is required.
Name	The IP address name. This field is required.
MAC address	The IP address MAC address. This field is optional.
Class name	The IP address class name. This field is optional.
Class parameters	The IP address-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Space name	The IP address space name. At the bottom of the list of columns of the CSV file, the existing spaces are also listed, select the space where you want import the address(es). This field is required.

8. Click on **[NEXT]**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **[NEXT]**. The page **CSV import parameters** opens:

- a. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
- b. Tick this box **Trigger the execution of all advanced properties** if you want to force the advanced property mechanisms during the import. These mechanisms depend on the configuration set on the page *Class parameters* and/or on the properties inherited from higher level(s), they can include properties impacting the IPAM, DNS, DHCP and Device Manager. For more details, refer to the section [IP Address Advanced Properties](#) or to the chapter [Managing the Interaction with the IPAM](#).

After the import, you can still trigger the execution of the configured mechanisms. Tick the IP addresses of your choice and in the menu select **Tools > Expert > Initialize rules**.

10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All addresses**. The IP addresses are now listed.

Importing Data to the DHCP

Within the DHCP module, you can import data on the scopes, ranges and statics pages in IPv4 and IPv6. The table below details where you can import them within the module.

Table 9.12. DHCP pages where you can import CSV files

DHCP page	Objects that can be imported	Option name in the menu Import
All scopes	Scopes	CSV scopes
	Ranges	CSV ranges
	Statics	CSV statics
All scopes (v6)	IPv6 scopes	CSV scopes (v6)
	IPv6 ranges	CSV ranges (v6)
	IPv6 statics	CSV statics (v6)
All ranges	Ranges	CSV ranges
All ranges (v6)	IPv6 ranges	CSV ranges (v6)
All statics	Statics	CSV statics
All statics (v6)	IPv6 statics	CSV statics (v6)

Before importing, keep in mind that:

- The advanced properties can automate the creation of entries in your database after the import. If all the advanced properties are activated, importing DHCP data may automatically update

the IPAM and DNS databases. If you do not want your import to impact other modules, edit the *Internal module setup* before importing DHCP data. For more details, refer to the chapter [Managing Advanced Properties](#).

To import an ISC, Alcatel-Lucent VitalQIP, Microsoft, Infoblox, MetalP or Nortel NetID configuration, refer to the chapter [Importing DHCP Data](#).

Importing Scopes

You can import several scopes coming from different DHCP configurations into the same server. If you plan on importing scopes into different servers, make sure that your CSV file contains a column dedicated to the server name.

Importing IPv4 Scopes

When importing a file containing IPv4 scope(s), on the page **CSV fields association**, the fields **Start address**, **DHCP server** and one of the fields specifying the scope size are required: **End address** or **Size** must be specified.

To import IPv4 scopes through a CSV file

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Click on the **Name** of the DHCP server or smart architecture of your choice. The page **All scopes** of the server opens.
3. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
4. In the menu, select **Import > CSV scopes**. The wizard **Import a CSV file** opens.
5. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
6. Click on **[NEXT]**. The page **CSV fields association** opens.
7. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
8. Select the columns of your CSV file you want to import.

The fields **Address**, **DHCP server** and one field specifying the scope size (**Prefix**, **Netmask** or **Size**) are required. All fields are detailed in the table below.

Table 9.13. DHCP scope import parameters

Parameter	Description
Name	The scope name. This field is optional.
Address	The scope first address. This field is required.
Prefix	The drop-down lists that specify the size of the scope(s), the number of IP addresses being imported. You must fill at least one of them.
Netmask	
Size	
Scope space	The scope space in the IPAM. This field is optional.
Shared network	The scope shared network. This field is optional.
Failover	The scope failover. This field is optional.
Class name	The scope class name. This field is optional.

Parameter	Description
Class parameters	The scope-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP server	The scope server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the scope(s). This field is required.

9. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
10. Click on **NEXT**. The page **DHCP options** opens. All the fields are optional, choose the data you want to import. For more details, refer to the chapter [Configuring DHCP Options](#).
11. Click on **NEXT**. The page **CSV import parameters** opens.
12. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
13. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
15. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 17.
16. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
17. Click on **CLOSE** to go back to the page **All scopes**. The scopes are now listed.

Importing IPv6 Scopes

When importing a file containing IPv6 scope(s), on the page **CSV fields association**, the fields **Start address**, **DHCP server** and one of the fields specifying the scope size are required: **End address** or **Size** must be specified.

To import IPv6 scopes through a CSV file

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Click on the **Name** of the DHCP server or smart architecture of your choice. The page **All scopes** of the server opens.
3. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
4. In the menu, select **Import > CSV scopes (v6)**. The wizard **Import a CSV file** opens.
5. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
6. Click on **NEXT**. The page **CSV fields association** opens.

7. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
8. Select the columns of your CSV file you want to import.

The fields **Start address**, **DHCP server**, any field specifying the scope size (**End address** or **Prefix**) are required. All fields are detailed in the table below.

Table 9.14. DHCPv6 scope import parameters

Parameter	Description
Name	The scope name. This field is optional.
Start address	The scope start address. This field is required.
End address	The drop-down lists that specify the size of the scope(s), the number of IP addresses being imported. You must fill at least one of them.
Prefix	
Scope space	The scope space in the IPAM. This field is optional.
Class name	The scope class name. This field is optional.
Class parameters	The scope-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP6 server	The scope server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the scope(s). This field is required.

9. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
10. Click on [NEXT](#). The page **DHCP options** opens. All the fields are optional, choose the data you want to import. For more details, refer to the chapter [Configuring DHCP Options](#).
11. Click on [NEXT](#). The page **CSV import parameters** opens.
12. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
13. Click on [CHECK](#). The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the validity report in the specified file format.
15. Click on [OK](#) to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 17.
16. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the import report in the specified file format.
17. Click on [CLOSE](#) to go back to the page **All scopes**. The scopes are now listed.

Importing Ranges

From the page *All ranges*, you can import IPv4 and IPv6 ranges. These ranges must be imported within an existing scope.

Importing IPv4 Ranges

When importing a file containing IPv4 range(s), on the page **CSV fields association**, the fields **Start address**, **DHCP server** and one of the fields specifying the scope size are required: **End address** or **Size** must be specified.

To import IPv4 ranges through a CSV file

1. In the sidebar, go to [DHCP > Ranges](#). The page **All ranges** opens.
2. On the right-end side of the menu, click on [v4](#). The page refreshes and the button turns black.
3. In the menu, select **Import > CSV ranges**. The wizard **Import a CSV file** opens.
4. Click on [BROWSE](#) to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on [NEXT](#). The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **Start address**, **DHCP server** and any field specifying the range size (**End address** or **Size**) are required. All fields are detailed in the table below.

Table 9.15. DHCP range import parameters

Parameter	Description
Start address	The range(s) start address. This field is required.
End address	The drop-down lists that specify the size of the range(s), the number of IP addresses being imported. You must fill at least one of them.
Size	
Failover channel	The range failover channel. This field is optional.
ACL	The range ACL. This field is optional.
Class name	The range class name. This field is optional.
Class parameters	The range-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP server	The range server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the range(s). This field is required.

8. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on [NEXT](#). The page **DHCP options** opens. All the fields are optional, choose the data you want to import. For more details, refer to the chapter [Configuring DHCP Options](#).

10. Click on **NEXT**. The page **CSV import parameters** opens.
11. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
12. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
14. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 16.
15. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
16. Click on **CLOSE** to go back to the page **All ranges**. The ranges are now listed.

Importing IPv6 Ranges

Note that there are no DHCP options for DHCPv6 ranges, so you cannot import them.

When importing a file containing IPv6 range(s), on the page **CSV fields association**, the fields **Start address**, **End address** and **DHCP6 server** are required.

To import IPv6 ranges through a CSV file

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV ranges (v6)**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The field **Start address**, **End address** and **DHCP6 server** are required. All fields are detailed in the table below.

Table 9.16. DHCPv6 range import parameters

Parameter	Description
Start address	The range start address. This field is required.
End address	The range last address. This field is required.
Class name	The range class name. This field is optional.

Parameter	Description
Class parameters	The range-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP6 server	The range server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the range(s). This field is required.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
13. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
15. Click on **CLOSE** to go back to the page **All ranges**. The ranges are now listed.

Importing Statics

From the page *All statics*, you can import IPv4 and IPv6 static reservations. These statics can be imported in a DHCP server or group.

If you are importing statics with IP address in a DHCP server, keep in mind that they are managed like leases by the server. For more details, refer to the section [Adding DHCPv4 Statics](#).

Importing IPv4 Statics

When importing a file containing IPv4 static reservation(s), on the page **CSV fields association**, the fields **DHCP server**, **MAC address** and **Client DUID** are required.

To import IPv4 statics through a CSV file

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV statics**. The wizard **Import a CSV file** opens.

4. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **[NEXT]**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **MAC address** and **DHCP server** are required. All fields are detailed in the table below.

Table 9.17. DHCP static import parameters

Field	Description
DHCP static name	The static name. This field is optional. For EfficientIP DHCP servers, if you specify the static name and leave the field <i>Option host-name</i> empty, the value of the DHCP option is used as the static name as well. You should specify either the <i>DHCP static name</i> or the <i>Option host-name</i> .
MAC address	The static MAC address. This field is required.
DHCP static IP address	The static IP address. This field is optional.
DHCP group	The static group. This field is optional.
DHCP static class name	The static class name. This field is optional.
Class parameters	The static-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP server	The static server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the static(s). This field is required.

8. Click on **[NEXT]**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **[NEXT]**. The page **DHCP options** opens. All the fields are optional, choose the data you want to import. For more details, refer to the chapter [Configuring DHCP Options](#).

For EfficientIP DHCP servers, you can specify the **Option host-name** and leave the field **DHCP static name** empty to use the value of the option as the name of the static. You should specify either the *DHCP static name* or the *Option host-name*.

10. Click on **[NEXT]**. The page **CSV import parameters** opens.
11. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
12. Click on **[CHECK]**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **[TEXT]**, **[HTML]**, or **[EXCEL]** to download the validity report in the specified file format.
14. Click on **[OK]** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If

you want to download that final report, refer to the next step. If you do not want to download it, go to step 16.

15. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
16. Click on **CLOSE** to go back to the page **All statics**. The static reservations are now listed.

Importing IPv6 Statics

When importing a file containing IPv6 static reservation(s), on the page **CSV fields association**, the fields **DHCP server**, **MAC address** and **Client DUID** are required.

To import IPv6 statics through a CSV file

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. In the menu, select **Import > CSV statics (v6)**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **NEXT**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **DHCP static name** and **DHCP server** are required. All fields are detailed in the table below.

Table 9.18. DHCPv6 static import parameters

Field	Description
DHCP static name	The static name. This field is required.
Static IP address	The static IP address. This field is optional.
MAC address	The drop-down lists that identify the static(s) client. It is required to fill at least one of them.
Client DUID	
DHCP group	The static group. This field is optional.
DHCP static class name	The static class name. This field is optional.
Class parameters	The static-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
DHCP6 server	The static server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the static(s). This field is required.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **DHCP options** opens. All the fields are optional, choose the data you want to import. For more details, refer to the chapter [Configuring DHCP Options](#).

10. Click on **NEXT**. The page **CSV import parameters** opens.
11. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
12. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
14. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 16.
15. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
16. Click on **CLOSE** to go back to the page **All statics**. The static reservations are now listed.

Importing Data to the DNS

Within the DNS module, you can import zones and resource records. The table below details where you can import them within the module.

Table 9.19. DNS pages where you can import CSV files

DNS page	Objects that can be imported	Option name in the menu ↶ Import
All zones	Zones	CSV zones
All RRs	Resource records	CSV RRs

Before importing, keep in mind that:

- The advanced properties can automate the creation of entries in your database after the import. If all the advanced properties are activated, importing DNS data may automatically update the IPAM and DHCP databases. If you do not want your import to impact other modules, edit the *Internal module setup* before importing DNS data. For more details, refer to the chapter [Managing Advanced Properties](#).

To import zones from a BIND or VitalQIP archive file, refer to the chapter [Importing DNS Data](#).

Importing Zones

When importing a file containing zone(s), on the page **CSV fields association**, the fields **DNS zone name**, **DNS zone type** and **DNS server name** are required.

To import zones through a CSV file

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the **Name** of the server of your choice. The page **All zones** of the server opens.
3. In the menu, select **↶ Import > CSV zones**. The wizard **Import a CSV file** opens.
4. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.

5. Click on **NEXT**. The page **CSV fields association** page appears.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

The fields **DNS zone name**, **DNS zone type** and **DNS server name** are required. All fields are detailed in the table below.

Table 9.20. Zone import parameters

Parameter	Description
DNS Zone name	The zone(s) name. This field is required.
DNS Zone type	The static type. This field is required.
Master IP address	The static master server IP address. This field is required when importing slave zone(s).
Forwarder IP address	The static forwarding server IP address. This field is optional.
DNS view	The static view name. This field is optional.
DNS server name	The static server name. At the bottom of the list of columns of the CSV file, the existing servers are also listed, select the server where you want to import the zone(s). This field is required.

8. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on **NEXT**. The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
13. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
15. Click on **CLOSE** to go back to the page **All zones**. The zones are now listed.

Importing Resource Records

When importing a file containing resource record(s), on the page **CSV fields association**, the fields **RR name**, **Value 1**, **Zone name** and **RR type** are required.

To import resource records through a CSV file

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the menu, select **Import > CSV RRs**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** page appears.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

The fields **RR name**, **Value 1** and **RR type** are required. All fields are detailed in the table below.

Table 9.21. Resource record import parameters

Parameter	Description
RR name	The RR(s) name. This name can be FQDN if you also import the column containing the <i>Zone name</i> . This field is required.
TTL	The RR TTL. This field is optional.
Value 1	The first piece of information in the column <i>Value</i> . This field is required.
Value 2	The extra information that can contain the column <i>Value</i> depending on the RR type. These fields are optional.
Value 3	
Value 4	
Value 5	
Value 6	
Value 7	
Zone name	
DNS view	The RR view name. This field is optional.
DNS server	The RR server name. This field is optional.
RR type	The RR type. At the bottom of the list of columns of the CSV file, the existing RR types are also listed, you can choose one of them. This field is required.

7. Click on **NEXT**. The page **CSV import parameters** opens.
8. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
9. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 11.
10. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
11. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it go to step 13.

12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
13. Click on **CLOSE** to go back to the page **All RRs**. The records are now listed.

Importing Data to NetChange

Within the module NetChange, you can import network devices with an IPv4 address and that manage interfaces with IPv4 or IPv6 addresses. The table below details where you can import them within the module.

Table 9.22. NetChange pages where you can import CSV files

NetChange page	Objects that can be imported	Option name in the menu ⏪ Import
All network devices	Network devices	CSV file

Importing Network Devices

When importing network device(s), only the field **Target space** is required.

To import network devices through a CSV file

1. In the sidebar, go to **NetChange > Network devices**. The page **All network devices** opens.
2. In the menu, select **⏪ Import > CSV network devices**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** page appears.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import as described in the table below:

Table 9.23. Network device import parameters

Parameter	Description
Address	The network device(s) IP address. This field is mandatory.
Community	The SNMPv1/v2c community string that would act as a password to access SNMP agent on the device(s). For SNMPv3, one or several SNMP profiles can be specified on the last page of the wizard. This field is optional.
SNMP version	The version of the SNMP protocol used on the network device(s). It must be either 1 for SNMPv1 or 2 for SNMPv2c. For SNMPv3, one or several SNMP profiles can be specified on the last page of the wizard. This field is optional.
Class	The network device(s) class. This field is optional.
Target space	The network device(s) space in the IPAM, i.e. the space that should list the IP address of the discovered items on the network device(s). This field is required. If you have classes enabled, it will appear on the next page of the wizard.

7. Tick the box **Expert mode** to specify more details regarding the device(s) information retrieval. If you have classes enabled, it will appear on the next page of the wizard. Edit the parameters according to your needs following the table below:

Table 9.24. SNMP parameters

Field	Description
SNMP port	The port that the SNMP service must use. By default, the port <i>161</i> is selected.
SNMP retries	Select the number of connection attempts when the server is in timeout. You can set it between <i>0</i> and <i>5</i> . By default, it is set to <i>2</i> attempts.
SNMP transfer timeout (minutes)	Set the number of minutes above which the SNMP transfer is aborted when you add or refresh a device. You can set it between <i>0</i> and <i>999</i> . By default, it is set to <i>0</i> .
Use bulk	If you use SNMPv2c or v3, you can choose to use a bulk transfer of data. This compact SNMP request method accelerates transfers by sending several requests at once. By default, it is set to <i>Yes</i> .
Use TCP	Choose to use the TCP protocol instead of the UDP when the network link is not reliable. By default <i>UDP</i> is used, the drop-down list is set to <i>No</i> .

- Click on **NEXT**. If you have classes enabled, the page **Class parameters** opens. You can set the class parameters you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters. The fields **Target space** and **Expert mode** are described in the previous step.
- Click on **NEXT**. The last page of the wizard appears. You can select the SNMP profile(s) to use in order to access the SNMP agent on the devices if they are not associated to a version or community string in the CSV file. **It is the only way to specify authentication parameters in SNMPv3.**

Table 9.25. SNMP profile information parameters

Parameter	Description
SNMP profiles configuration	The SNMP profiles available. By default there are 3 profiles, <i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i> , but you can create as many profiles as needed to display them in this list. For more details, refer to the section Managing SNMP Profiles . Select a profile and click on + to move it to the list <i>Selected profiles</i> .
Selected profiles	This field lists the SNMP profiles to use in order to retrieve the device(s) information. SOLIDserver tries all the profiles on the device(s), following the list order. To remove a profile from the list, select it and click on - .

If you do not select any, NetChange uses the profile *standard v2c*.

- Click on **OK** to complete the operation. The **Report** opens and work for a while: the import progression is visible. Once the import is over, the report lists the IP addresses imported as well as the existing ones. If you want to download that final report, refer to the next step. If you do not want to download it, go to step *12*.
- In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
- Click on **CLOSE** to go back to the page **All network devices**. The devices are now listed.

Importing Data to Device Manager

Within Device Manager module, you can import data from the page *All Devices* and *All ports & interfaces*. The table below details where you can import them within the module.

Table 9.26. Device Manager pages where you can import CSV files

Device Manager page	Objects that can be imported	Option name in the menu ↶ Import
All devices	Devices	CSV devices
	Ports and/or interfaces	CSV interfaces
All ports & interfaces	Ports and/or interfaces	CSV interfaces

Importing Devices

When importing a file containing device(s), on the page **CSV fields association**, only the field **Name** is required.

To import devices through a CSV file

1. In the sidebar, go to the **Device Manager > Devices**. The page **All devices** opens.
2. In the menu, select **↶ Import > CSV devices**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

Only the field **Name** is required. All fields are detailed in the table below.

Table 9.27. Device import parameters

Parameter	Description
Name	The device name. This field is required.
Class name	The device class name. This field is optional.
Class parameters	The device-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.

7. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **NEXT**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.

12. Click on **[OK]** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **[TEXT]**, **[HTML]**, or **[EXCEL]** to download the import report in the specified file format.
14. Click on **[CLOSE]** to go back to the page **All devices**. The devices are now listed.

Importing Ports & Interfaces

When importing a file containing ports and interface(s), on the page **CSV fields association**, only the fields **Name**, **Type** and **Device** are required.

To import ports and/or interfaces through a CSV file

1. In the sidebar, go to **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
2. In the menu, select **Import > CSV interfaces**. The wizard **Import a CSV file** opens.
3. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **[NEXT]**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

Only the field **Name** is required. All fields are detailed in the table below.

Table 9.28. Port and/or interface import parameters

Parameter	Description
Name	The port and/or interface name. This field is required.
Type	The port and/or interface type. This field is required.
MAC address	The port and/or interface MAC address. This field is optional.
Class name	The port and/or interface class name. This field is optional.
Class parameters	The port and/or interface-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Device	The port and/or interface device name. At the bottom of the list of columns of the CSV file, the existing devices are also listed, select the device where you want to import the port(s) and/or interface(s). This field is required.

7. Click on **[NEXT]**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **[NEXT]**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.

10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All ports & interfaces**. The ports and interfaces are now listed.

Importing Data to VLAN Manager

Within VLAN Manager module, you can import domains, ranges and VLANs. The table below details where you can import them within the module.

Table 9.29. VLAN Manager pages where you can import CSV files

VLAN Manager page	Objects that can be imported	Option name in the menu ⌂ - Import
All domains	Domains	CSV domains
	Ranges	CSV ranges
	VLANs	CSV VLANs
All ranges	Ranges	CSV ranges
	VLANs	CSV VLANs
All VLANs	VLANs	CSV VLANs

Importing VLAN Domains

When importing a file containing VLAN domain(s), on the page **CSV fields association**, only the fields **Name**, **Start ID** and **End ID** are required.

To import VLAN domains through a CSV file

1. In the sidebar, go to **⇒ VLAN Manager > Domains**. The page **All Domains** opens.
2. In the menu, select **⌂ - Import > CSV domains**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

The fields **Name**, **Start ID** and **End ID** are required. All fields are detailed in the table below.

Table 9.30. VLAN domain import parameters

Parameter	Description
Name	The domain name. This field is required.
Start ID	The domain first VLAN ID. This field is required.
End ID	The domain last VLAN ID. This field is required.
Description	The domain description. This field is optional.
Class name	The domain class name. This field is optional.
Class parameters	The domain-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.

7. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on [NEXT](#). The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on [CHECK](#). The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the validity report in the specified file format.
12. Click on [OK](#) to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the import report in the specified file format.
14. Click on [CLOSE](#) to go back to the page **All domains**. The VLAN domains are now listed.

Importing VLAN Ranges

When importing a file containing VLAN ranges(s), on the page **CSV fields association**, only the fields **Name**, **Start ID**, **End ID** and **Domain** are required.

To import VLAN ranges through a CSV file

1. In the sidebar, go to [VLAN Manager > Ranges](#). The page **All ranges** opens.
2. In the menu, select [Import > CSV ranges](#). The wizard **Import a CSV file** opens.
3. Click on [BROWSE](#) to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on [NEXT](#). The **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).

6. Select the columns of your CSV file you want to import.

The fields **Name**, **Start ID**, **End ID** and **Domain** are required. All fields are detailed in the table below.

Table 9.31. VLAN range import parameters

Parameter	Description
Name	The range name. This field is required.
No ID overlapping	The VLAN ID overlapping policy for the range. By default it is set to <i>yes</i> , the overlapping is disabled: the ranges can only contain unique VLAN IDs. This field is optional.
Start ID	The range first VLAN ID. This field is required.
End ID	The range last VLAN ID. This field is required.
Description	The range description. This field is optional.
Class name	The range class name. This field is optional.
Class parameters	The range-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.
Domain	The range domain name. This field is required.

7. Click on [NEXT](#). The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on [NEXT](#). The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on [CHECK](#). The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the validity report in the specified file format.
12. Click on [OK](#) to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on [TEXT](#), [HTML](#), or [EXCEL](#) to download the import report in the specified file format.
14. Click on [CLOSE](#) to go back to the page **All ranges**. The VLAN ranges are now listed.

Importing VLANs

When importing a file containing VLAN ranges(s), on the page **CSV fields association**, only the fields **VLAN ID**, **Range** and **Domain** are required.

To import VLANs through a CSV file

1. In the sidebar, go to [VLAN Manager > VLANs](#). The page **All VLANs** opens.
2. In the menu, select [Import > CSV VLANs](#). The wizard **Import a CSV file** opens.

3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

The fields **VLAN ID**, **Range** and **Domain** are required. All fields are detailed in the table below.

Table 9.32. VLAN import parameters

Parameter	Description
Name	The VLAN name. This field is optional.
VLAN ID	The VLAN ID. This field is required.
Range	The VLAN range name. This field is required.
Domain	The VLAN domain name. This field is required.

7. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **NEXT**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All VLANs**. The VLANs are now listed.

Importing Data to VRF

Within the VRF module, you can import VRFs and VRF route targets. The table below details where you can import them within the module.

Table 9.33. VRF pages where you can import CSV files

VRF page	Objects that can be imported	Option name in the menu ← Import
All VRFs	VRFs	CSV VRFs
	VRF Route Targets	CSV VRF Route Targets

VRF page	Objects that can be imported	Option name in the menu  Import
All VRF Route Targets	VRF Route Targets	CSV VRF Route Targets

Importing VRFs

When importing VRF(s), the fields **VRF name** and **VRF RD** are required.

To import VRFs through a CSV file

1. In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
2. In the menu, select  **Import > CSV VRFs**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

The fields **VRF name** and **VRF RD** are required. All fields are detailed in the table below.

Table 9.34. VRF import parameters

Parameter	Description
VRF name	The VRF name. This field is required.
VRF RD	The VRF RD. This field is required.
VRF comment	The VRF comment. This field is optional.
VRF class name	The VRF class name. This field is optional.
VRF class parameters	The VRF-related combination of parameters, in URL format. It can contain all the class parameters of the resource, in which case you do not need to specify them one by one on the page <i>Class parameters</i> . This field is optional.

7. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **NEXT**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.
11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.

13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All VRFs**. The VRFs are now listed.

Importing VRF Route Targets

You can import VRF Route Targets to set up communication between existing VRFs.

When importing VRF route target(s), the fields **Source RD of the VRF Route Targets** and **Target RD of the VRF Route Targets** are required. As the VRFs are already in the database, the name is retrieved and displayed on the page once the Route Targets are imported.

To import VRF Route Targets through a CSV file

1. In the sidebar, go to **VRF > VRF Route Targets**. The page **All VRF Route Targets** opens.
2. In the menu, select **Import > CSV VRF Route Targets**. The wizard **Import a CSV file** opens.
3. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
4. Click on **NEXT**. The page **CSV fields association** opens.
5. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
6. Select the columns of your CSV file you want to import.

The fields **Source RD of the VRF Route Targets** and **Target RD of the VRF Route Targets** are required. All fields are detailed in the table below.

Table 9.35. VRF route target import parameters

Parameter	Description
Source RD of the VRF Route Targets	The source VRF of the Route Target. This field is required.
Target RD of the VRF Route Targets	The target VRF of the Route Target. This field is required.
Imported VRF Route Target	The import VRF Route Target parameter. This field is optional.
Exported VRF Route Target	The export VRF Route Target parameter. This field is optional.

7. Click on **NEXT**. The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
8. Click on **NEXT**. The page **CSV import parameters** opens.
9. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
10. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 12.

11. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
12. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 14.
13. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
14. Click on **CLOSE** to go back to the page **All VRF Route Targets**. The VRF Route Targets are now listed.

Importing Data to SPX

When it comes to importing RIPE or APNIC objects, within the SPX module, you can only import SPX aut-nums. The table below details where you can import them within the module.

Table 9.36. SPX pages where you can import CSV files

SPX page	Objects that can be imported	Option name in the menu ⌂ Import
All AS Numbers	SPX aut-nums	SPX aut-nums

However, you can also import:

- SPX allocated and assigned networks from the IPAM page All networks, and
- SPX users, i.e. persons, from the Administration page All users.

Importing SPX Allocated Networks

Once your SPX configuration is complete, you can import the networks that the RIPE or APNIC allocated to you. For more details, refer to the chapter [Configuring SPX](#).

During the import, the option (box) *Use the "ripe.db.inetnum" file stored in the Local files listing* allows you to use the "ripe.db.inetnum" file if you uploaded it to the Local files listing before performing the import. It allows to work with the file content rather than connecting to the RIPE or APNIC using an Internet connection to obtain the assigned network details. For more details on how to upload a file to the Local files listing, refer to the section [Managing Files from the Local Files Listing](#).

Note that, following the IPAM hierarchy, your allocated network(s) must belong to a space.

To import an IPv4 SPX allocated network

1. In the sidebar, go to **⌂ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4**.
3. In the menu, select **⌂ Import > SPX allocated networks**. The wizard **Importing SPX allocated networks** opens.
4. You can tick the box **Use the "ripe.db.inetnum" file stored in the Local files listing** if you want.
5. In the drop-down list **Maintainer**, select the maintainer of your choice.

6. In the drop-down list **Target space**, select the space of your choice. If you are importing from the page *All networks* of a specific space, it is already selected.
7. In the drop-down list **PA allocated network class**, you can choose a class if you manage an allocated network of Provider Aggregatable IP addresses.
8. In the drop-down list **PI allocated network class**, you can choose a class if you manage an allocated network of Provider Independent IP addresses.
9. Click on to complete the operation. The report opens and closes, the page refreshes. The allocated network is listed.

To import an IPv6 SPX allocated networks

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on .
3. In the menu, select **↔ Import > SPX allocated networks (v6)**. The wizard **Importing IPv6 SPX allocated networks** opens.
4. You can tick the box **Use the "ripe.db.inetnum" file stored in the Local files listing** if you want.
5. In the drop-down list **Maintainer**, select the maintainer of your choice.
6. In the drop-down list **Target space**, select the space of your choice. If you are importing from the page *All networks* of a specific space, it is already selected.
7. In the drop-down list **PA allocated network class**, you can choose a class if you manage an allocated network of Provider Aggregatable IPv6 addresses.
8. In the drop-down list **PI allocated network class**, you can choose a class if you manage an allocated network of Provider Independent IPv6 addresses.
9. Click on to complete the operation. The report opens and closes, the page refreshes. The allocated network is listed.

Importing SPX Assigned Networks

Once your configuration with the SPX is complete and you have imported your allocated networks in a space, you can import your existing assigned networks if you have any. The RIPE or APNIC assigned networks correspond to the assigned networks in the IPAM hierarchy. For more details, refer to the chapter [Configuring SPX](#).

During the import, the option (box) *Use the "ripe.db.inetnum" file stored in the Local files listing* allows you to use the "ripe.db.inetnum" file if you uploaded it to the Local files listing before performing the import. It allows to work with the file content rather than connecting to the RIPE or APNIC using an Internet connection to obtain the assigned network details. For more details on how to upload a file to the Local files listing, refer to the section [Managing Files from the Local Files Listing](#).

Once you imported your network objects, editing the content of your assigned networks follows the same procedures as regular assigned networks. For more details, refer to the chapters [Managing Pools](#) and [Managing IP Addresses](#).

To import IPv4 SPX assigned networks

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on .

3. In the menu, select **⏪ Import > SPX assigned networks**. The wizard **Importing SPX assigned networks** opens.
4. You can tick the box **Use the "ripe.db.inetnum" file stored in the Local files listing** if you want.
5. In the drop-down list **Maintainer**, select the maintainer of your choice.
6. In the drop-down list **Destination space**, select the space of your choice. If you are importing from the page *All networks* of a specific space, it is already selected.
7. In the drop-down list **PA Assigned network class**, you can choose a class if you manage assigned networks of Provider Aggregatable IP addresses.
8. In the drop-down list **PI Assigned network class**, you can choose a class if you manage assigned networks of Provider Independent IP addresses.
9. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. The assigned networks are listed.

To import IPv6 SPX assigned networks

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v6**.
3. In the menu, select **⏪ Import > SPX assigned networks (v6)**. The wizard **Importing IPv6 SPX assigned networks** opens.
4. You can tick the box **Use the "ripe.db.inetnum" file stored in the Local files listing** if you want.
5. In the drop-down list **Maintainer**, select the maintainer of your choice.
6. In the drop-down list **Destination space**, select the space of your choice. If you are importing from the page *All networks* of a specific space, it is already selected.
7. In the drop-down list **PA Assigned network class**, you can choose a class if you manage assigned networks of Provider Aggregatable IPv6 addresses.
8. In the drop-down list **PI Assigned network class**, you can choose a class if you manage assigned networks of Provider Independent IPv6 addresses.
9. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. The assigned networks are listed.

Once you imported your assigned networks, you can edit them from the GUI. Any change is sent to the RIPE or APNIC using the update method that you selected during the maintainer configuration (post or email).

You can also add assigned networks from the GUI. These new objects are also communicated to the RIPE or APNIC. For more details, refer to the section [Adding SPX Assigned Networks](#).

Importing SPX Persons

You can import existing SPX persons from the module Administration, on the page Users. In the GUI, they are managed listed like the other users. The main goal of importing SPX persons is to edit them from the GUI, any change is sent to the RIPE or APNIC following the update method you selected when configuring the maintainer.

You can create a group for your SPX persons to gather them but, unlike standard users managed via the appliance, there is no need to grant them specific rights.

The SPX persons listed on the Users page do not have access to the appliance if you do not grant them rights (through the group they belong to) or configure credentials for them.

To import SPX persons

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. In the menu, select  **Import > SPX persons**. The wizard **Importing SPX persons (users)** opens.
4. In the drop-down list **Maintainer**, select the maintainer of your choice.
5. Click on to complete the operation. The report opens and closes, the page refreshes. The SPX persons are listed among the users.

Importing SPX Aut-nums and AS Policies

You can import Autonomous System numbers (aut-nums) on the AS Numbers page.

Importing of AS Numbers also imports AS routing policies. The routing policy is described by enumerating all neighboring AS Number with which routing information is exchanged, they are all listed in the page All policies. For each neighbor, the routing policy is described in terms of exactly what is being sent (announced) and allowed (accepted). That way, each aut-num contains policies that describes what can be implemented and enforced locally by said AS Number.

Keep in mind the page All policies is accessible from the page All AS Numbers. You can access it through the breadcrumb.

To import SPX aut-nums

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. In the menu, select  **Import > SPX aut-nums**. The wizard **Importing SPX aut-nums (AS Numbers)** opens.
3. In the drop-down list **Maintainer**, select the maintainer of your choice.
4. In the drop-down list **Class name**, you can select a class to apply to the aut-nums you are importing.
5. Click on to complete the operation. The report opens and closes, the page refreshes. The aut-nums are listed.
6. In the column **AutNum name**, click on the name of the aut-num of your choice. The page **All policies** displays the policies of this AS Number.

Importing Data to the Administration Module

Within the Administration module, you can import data on the Groups, Users and Custom data pages. The table below details where you can import them within the module.

Table 9.37. Administration module pages where you can import CSV files

Administration page	Objects that can be imported	Option name in the menu  Import
Groups	Groups of users	CSV groups
Users	Users	CSV file

Administration page	Objects that can be imported	Option name in the menu  Import
	SPX users ^a	SPX persons
Custom data	Custom data	CSV custom data

^aTo import RIPE or APNIC users, i.e. persons, refer to the section [Importing SPX Persons](#).

Importing Groups of Users

When importing a file containing group(s) of users, on the page **CSV fields association**, only the field **Name** is required.

To import groups through a CSV file

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. In the menu, select  **Import > CSV groups**. The wizard **Import a CSV file** opens.
4. Click on to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on . The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

Only the field **Name** is required. All fields are detailed in the table below.

Table 9.38. Group of users import parameters

Parameter	Description
Name	The group of users name. This field is required.
Description	The group of users description. This field is optional.
Category	The group of users category. This field is optional.
Class name	The group of users class name. This field is optional.
Group parent name	The name of the group of users' parent group. This allows to copy the rights of the selected group and apply them to the group(s) your are importing. This field is optional.

8. Click on . The page **Class parameters** opens. If you want to import class parameters and did not import all your class parameters on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional, they are named after the parameters.
9. Click on . The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on . The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on , , or to download the validity report in the specified file format.

13. Click on **[OK]** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **[TEXT]**, **[HTML]**, or **[EXCEL]** to download the import report in the specified file format.
15. Click on **[CLOSE]** to go back to the page **Group**. The groups are now listed.

Importing Users

To import RIPE or APNIC users, i.e. persons, refer to the section [Importing SPX Persons](#).

When importing a file containing user(s), on the page **CSV fields association**, only the field **Login** is required.

To import users through a CSV file

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. In the menu, select **⬅ Import > CSV file**. The wizard **Import a CSV file** opens.
4. Click on **[BROWSE]** to select the CSV file to import. The selected file is visible in the field **File name**.
5. Click on **[NEXT]**. The page **CSV fields association** opens.
6. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
7. Select the columns of your CSV file you want to import.

Only the field **Name** is required. All fields are detailed in the table below.

Table 9.39. User import parameters

Parameter	Description
Login	The user login. This field is required.
First name	The user first name. This field is optional.
Last name	The user last name. This field is optional.
Email	The user email address. This field is optional.
Password	The password the user should use to access SOLIDserver. This field is optional.
Description	The user description. This field is optional.
Authentication method	The user authentication method. This field is optional.
Default page	The user default page. This field is optional.
Class name	The user class name. This field is optional.
Maintainer group	The user maintainer group. This field is optional.

8. Click on **[NEXT]**. The page **Class parameters**. If you want to import class parameters for the class you selected on the previous page, you can set the ones you need one by one in the drop-down lists. All the fields are optional.
9. Click on **[NEXT]**. The page **CSV import parameters** opens.

10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.
12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
13. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
15. Click on **CLOSE** to go back to the page **Users**. The users are now listed.

Importing Custom Data

Within a custom DB you can import one or several custom data entries. When importing a file containing custom data, on the page **CSV fields association**, only the field **Value 1** is required. All the fields are named *Value <number>* to match the default column name of a custom DB.

To import custom data through a CSV file

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the column **Name**, click on the name of the custom database of your choice. The page **Custom data** of that database opens.
4. In the menu, select **Import > CSV custom data**. The wizard **Import a CSV file** opens.
5. Click on **BROWSE** to select the CSV file to import. The selected file is visible in the field **File name**.
6. Click on **NEXT**. The page **CSV fields association** opens.
7. Specify the format of the import file using the fields **Delimiter**, **Enclosure**, **Input format**, **Skip the first line** and set in as a **Template** if you want. For more details, refer to the table [CSV file basic parameters](#).
8. Select the columns of your CSV file you want to import.

Only the field **Value 1** is required. There are in total 10 fields named **Value 1** through to **Value 10**.

9. Click on **NEXT**. The page **CSV import parameters** opens.
10. In the drop-down list **Existing records**, select either *Replace* to overwrite the existing records that have the same name or *Don't replace* to add the items to the listing. *Don't replace* is selected by default.
11. Click on **CHECK**. The page **Check the validity of the CSV file** opens and displays a report indicating the total amount of correct lines within the file. If you want to download the validity report, refer to the next step. If you do not want to download it, go to step 13.

12. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the validity report in the specified file format.
13. Click on **OK** to accept the validity check report results. The page **Import data from a CSV file** opens and displays a report indicating the total number of spaces actually imported. If you want to download that final report, refer to the next step. If you do not want to download it, go to step 15.
14. In the section **Export format** of the wizard, click on **TEXT**, **HTML**, or **EXCEL** to download the import report in the specified file format.
15. Click on **CLOSE** to go back to the page **Custom data**. The entries are now listed.

Managing Import Templates

Every import template you created is listed on the page **Import/Export Templates Management** of each module. On this page, you can either rename or delete the templates.

This page is composed of as many panels as there are pages where you can import data in the module. Each panel lists all the templates configured on the page, the import templates are listed as follows: *Import: <template_name>*.

To rename an import template

1. In the sidebar, go to the module of your choice.
2. In the menu, select **⋮** > **Extra options** > **Import/Export templates management**. The page **Import/Export Templates Management** opens.
3. In the panel of your choice, select the *Import: <template_name>* you want to rename.
4. Click on **RENAME**. The wizard **Rename template** opens.
5. In the field **New Name**, rename your template.
6. Click on **OK** to complete the operation. The report opens and closes. The name changes in the list.

To delete an import template

1. In the sidebar, go to the module of your choice.
2. In the menu, select **⋮** > **Extra options** > **Import/Export templates management**. The page **Import/Export Templates Management** opens.
3. In the panel of your choice, select the *Import: <template_name>* you want to delete.
4. Click on **DELETE**. The wizard **Delete template** opens.
5. In the field **New Name**, rename your template.
6. Click on **OK** to complete the operation. The report opens and closes. The template is no longer listed.

Chapter 10. Importing IPAM Data

SOLIDserver offers several ways of importing existing IP addresses organizations without having to recreate them manually in the GUI.

You can import:

- [VitalQIP data](#).
- [Nortel NetID data](#).

Before importing, keep in mind that:

- The advanced properties can automate the creation of entries in your database after the import. If all the advanced properties are activated, importing IPAM data may automatically update the DHCP and DNS databases. If you do not want your import to impact other modules, edit the *Internal module setup* before importing IPAM data. For more details, refer to the chapter [Managing Advanced Properties](#).
- Other advanced options, like the IPv4 to IPv6 transition, the IPAM / Device Manager interaction properties and the IPAM/VLAN interaction properties, may also update your database automatically. For more details, refer to the chapters [Setting Up a Transition From IPv4 to IPv6](#), [Managing the Interaction with the IPAM](#) and [Managing the IPAM/VLAN Interaction](#).

To import spaces, networks, pools or IP addresses from a CSV file, refer to the chapter [Importing Data from a CSV File](#).

Importing a VitalQIP Export

From the page *All spaces* you can import VitalQIP data in a *.qef* file. This file includes networks and addresses so you need to import only one file to manage all your organization from the GUI.

To import VitalQIP data into SOLIDserver:

- The QIP export file (**.qef*) must belong to an archive file which extension is *.tar*, *.gz*, *.tgz*, *.tar.gz* or *.zip*.
- The *.qef* file must be located at the root of the archive file.
- The archive file does not need to include the files **_aud.qef* as they are not relevant to the import and might make your import take longer.

To import Vital QIP data

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. In the menu, select **Import > QIP data**. The wizard **Import entries from file** opens.
3. Click on **BROWSE** to search for the CSV file to import. A window opens to help you browse through folders, select the needed file.
4. Click on **Open**. The window closes and the file is visible in the field **File name**.
5. Click on **NEXT**. The page **Select a space** opens.
6. In the drop-down list **Space**, select one of the following options:

Table 10.1. Space field available options

Field	Description
Create space per organization	Select this option to create a space for each of the organizations that you are importing. This option is selected by default.
Existing spaces	Select a space to import your data to be imported in one of your existing spaces.

Note that all the data the space contains is imported as well (block-type networks, subnet-type networks and addresses).

- In the drop-down list **Network (block) class**, select an existing class¹ to be applied to the block-type networks you are importing. If no class exists or is enabled, only *None* is listed.
- In the drop-down list **Network (subnet) class**, select an existing class to be applied to the subnet-type networks you are importing. If no class exists or is enabled, only *None* is listed.
- Click on to complete the operation. The report opens and closes. The data is listed according to your import configuration.

Importing Nortel NetID IP Address Data

You can import NetID networks, subnet-type networks and host addresses. The Nortel NetID database must be exported as a text file with a comma or semi-colon data delimiter. Here below, are listed the Nortel NetID's fields that SOLIDserver can import via a CSV files.

To avoid missing any parameters or losing any data, we recommend that you follow the module hierarchy during these imports in an existing space: block-type networks, subnet-type networks and finally IP addresses .

For more details regarding CSV imports, refer to the section [Importing Data to the IPAM](#).

Importing Nortel NetID Networks

In SOLIDserver, the Nortel NetID *Networks* are imported as **networks (blocks)**. Here below are listed the fields equivalence between the two appliances to help you go through with the networks import.

Table 10.2. Nortel NetID network fields name when importing networks

Nortel NetID field	SOLIDserver field
Network number	First address
Network name	Name
Subnet type	-
CIDR mask	-
Subnet mask	Netmask

For more details regarding CSV imports, refer to the section [Importing Data to the IPAM](#).

Importing Nortel NetID Subnets

In SOLIDserver, the Nortel NetID *Subnets* are also called **networks (subnets)**. However, the fields to describe them differ. For more details, refer to the table below.

¹The classes listed were created and enabled on the page Class Studio and apply to the IPAM networks. For more details, refer to the chapter [Configuring Classes](#).

Table 10.3. Nortel NetID subnet fields name when importing networks

Nortel NetID field	SOLIDserver field
Network number	Address
Network name	Name
Subnet type	-
CIDR mask	-
Subnet mask	Netmask

For more details regarding CSV imports, refer to the section [Importing Data to the IPAM](#).

Importing Nortel NetID Host Addresses

In SOLIDserver, the Nortel NetID *Host addresses* are also called **addresses**. Here below are listed the fields equivalence between the two appliances to help you go through with the networks import.

Table 10.4. Nortel NetID host addresses fields name when importing networks

Nortel NetID field	SOLIDserver field
Host address	IP address
Domain name	Domain
Client ID	-
MAC address	MAC address
CIMAC type	-
Custom fields	-

For more details regarding CSV imports, refer to the section [Importing Data to the IPAM](#).

Chapter 11. Importing DHCP Data

SOLIDserver offers several ways of importing data from legacy DHCP to all other DHCP servers managed from the GUI, these wizards are all the more useful during a migration. You can import:

- [ISC configuration files](#).
- [Alcatel-Lucent VitalQIP configuration files](#).
- [Microsoft configuration files](#).
- [Infoblox configuration files](#).
- [MetalP configuration files](#).
- [Nortel NetID configuration files](#).

Before importing, keep in mind that:

To import DHCP or DHCPv6 scopes, ranges and statics, refer to the section [Importing Data to the DHCP](#).

Importing an ISC DHCP Configuration

You can import DHCP configurations from the ISC DHCP software in IPv4. Through this import, the whole DHCP server configuration create the following elements within an EfficientIP DHCP server: scopes, ranges, leases, statics, groups and DHCP options. However, there are some restrictions:

- **Scopes restriction:** if the server you are importing contains overlapping scopes, only the first scope are imported, the rest is ignored.
- **Statics restriction:** statics associated to an IP address not included in one the scopes you are importing are ignored.
- **Shared network restriction:** shared network options are ignored.
- **DHCP options restriction:** only standard options are supported during the import. If the server was configured using *non standard* DHCP options, they can only be imported if they were previously defined either in the configuration file or within the SOLIDserver appliance. However, you can configure conditional options afterward using the DHCP ACLs.

The ISC DHCP loads its configuration from the file named *dhcpd.conf*. This file contains the whole configuration of the DHCP server. You can import this file directly from SOLIDserver GUI at scope level.

EfficientIP recommends reducing all lease times to one hour before switching to the new DHCP server in order to minimize the risk of duplicating IP address assignments during the transition from the legacy DHCP server to SOLIDserver. This measure ensures that when you turn off your legacy DHCP servers, the DHCP clients quickly move to SOLIDserver when their lease renewal efforts fail: they broadcast their first DISCOVER message and get an answer within the hour.

To import an ISC DHCP configuration

1. In the sidebar, go to [DHCP > Scopes](#). The page **All scopes** opens.

2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **← Import > ISC DHCP**. The **Import an ISC dhcpd.conf file** wizard opens.
4. Click on **[BROWSE]** to find the ISC *dhcpd.conf* file. Once you clicked on Open, the file is visible in the wizard **File name** field.
5. In the drop-down list **DHCP server**, select the target server.
6. Click on **[OK]** to complete the operation. The report opens and closes. The file is listed.

This procedure also works from the page All scopes of the server for which you want to import the ISC configuration: at server level click on the name of the server concerned, once on the page All scopes, follow the procedure from step 4. The server is automatically selected in the drop-down list DHCP server.

Several ISC configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted: if the configurations conflict with each other all the different elements are added. In the same way, if two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

Importing an Alcatel-Lucent VitalQIP Configuration

You can import DHCP configurations coming from Alcatel-Lucent VitalQIP solution in IPv4. Before importing a VitalQIP configuration, note that:

- The DHCP range concept does not exist in VitalQIP, each address identified as an object in VitalQIP can become a dynamic assignment.
- VitalQIP dynamic objects are imported as a DHCP range.
- If several VitalQIP contiguous dynamic objects are imported, only one dynamic DHCP range is added if all all dynamic objects share the same DHCP option set.
 - Only one dynamic DHCP range is added if all all dynamic objects share the same DHCP option set.
 - Several DHCP ranges are added if the dynamic objects do not share the same options
- The entire configuration file of a VitalQIP DHCP loads is described in the file *dhcpd.conf* that can be imported directly in the GUI on the page All scopes of a specific server.
- Several VitalQIP configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted: if the configurations conflict with each other all the different elements are added. In the same way, if two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

SOLIDserver supports the following VitalQIP configuration of DHCP:

- DHCP server options.
- Scopes.
- Scope options.
- Ranges.
- Range options.

- Address pools.
- Static reservations.
- Static reservations options.
- DHCP option definitions.

To properly import a VitalQIP configuration file, you must:

1. Follow the procedure [To import a VitalQIP DHCP configuration](#).
2. Aggregate the static and range options to apply them at scope level, as detailed in the section [Aggregating DHCP Options from Ranges or Statics](#).

EfficientIP recommends reducing all lease times to one hour before switching to the new DHCP server in order to minimize the risk of duplicating IP address assignments during the transition from the legacy DHCP server to SOLIDserver. This measure ensures that when you turn off your legacy DHCP servers, the DHCP clients quickly move to SOLIDserver when their lease renewal efforts fail: they broadcast their first DISCOVER message and get an answer within the hour.

To import a VitalQIP DHCP configuration

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **Import > QIP DHCP**. The **Import a QIP DHCP configuration file** wizard opens.
4. Click on **[BROWSE]** to find the VitalQIP dhcpd.conf file. Once you clicked on Open, the file is displayed in the field **File name**.
5. In the drop-down list **DHCP server**, select the target server.
6. Click on **[OK]** to complete the operation. The report opens and closes.

Aggregating DHCP Options from Ranges or Statics

If you imported an external DHCP configuration and/or a specific DHCP option is configured across all the ranges or statics of a scope, you can aggregate it, is apply it, on the scope managing the ranges or statics.

This aggregation automates a homogeneous configuration of DHCP options on each of the scopes of a specific server. Keep in mind that it:

1. Analyzes the DHCP options configured on the ranges or statics with IP address of a scope.
2. If one DHCP option is configured and has the same value on all the ranges or statics with IP address of a scope:
 - a. The DHCP option is applied to the parent scope.
 - b. The DHCP option is deleted at range or static level, as it is automatically propagated from the scope down.
3. If one DHCP option is configured on all the ranges or statics with IP address of a scope, but their value is not the same on all the objects, the option is not aggregated at scope level.

To aggregate range options

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Tick the server(s) of your choice.
3. In the menu, select **Edit > Aggregate range options**. The wizard opens.
4. Click on **OK** to complete the operation. The report opens and closes.

To aggregate static options

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Tick the server(s) of your choice.
3. In the menu, select **Edit > Aggregate static options**. The wizard opens.
4. Click on **OK** to complete the operation. The report opens and closes.

Importing a Microsoft DHCP Configuration

You can import DHCP configurations from Microsoft in IPv4 from one file generated by the Microsoft *netsh* command on your DHCP servers. The configuration files can be imported from Microsoft Windows Servers 2008, 2008 R2, 2012 R2 or 2016.

Microsoft *netsh* commands for DHCP offer a command-line tool that helps administrating the DHCP servers and provides an equivalent alternative to console-based management. You can run these commands from the command prompt of a Windows Server or from the command prompt for the *netsh* DHCP context.

To run netsh commands from the command prompt of a Windows Server, you must include "netsh dhcp" before every command and parameter as detailed below:

```
C:\netsh dhcp server \\myservername dump > C:\dump_dhcp.txt
```

This command, generated from your Microsoft DHCP server, allows to import its whole configuration including:

- Definition of DHCP server options.
- Scopes.
- Scope options.
- Ranges.
- Address pools.
- Reservations.
- Reservation options.
- Exclusions.

Because several Microsoft DHCP configuration files can be imported into the same EfficientIP DHCP server, DHCP options are not imported at the server level. They must be manually configured.

Microsoft allows creating only one range per network (i.e. scope) and then exclude the ranges of IP addresses you do not need. Unlike Microsoft, from SOLIDserver you can configure several ranges in one scope. On the page *All ranges*, the list of Microsoft imported ranges does not display

exclusion ranges, it only lists the ranges created. For more details regarding EfficientIP and Microsoft DHCP management differences, refer to the explanation [SOLIDserver DHCP configuration vs. Microsoft DHCP configuration](#).

In the same way, when a Microsoft DHCP range contains a reservation, EfficientIP imports a reservation wrapped around two DHCP ranges.

Keep in mind that:

- **With Win2008R2 it is impossible to create a static outside a range.**
- You cannot import failover relationships. To manage Microsoft failover relationships, refer to the section [Managing Agentless Microsoft DHCP Servers](#).
- After importing a Microsoft configuration, you may need to aggregate the DHCP options set at range or static level to apply from the scope down. For more details, refer to the section [Aggregating DHCP Options from Ranges or Statics](#).

To import a Microsoft DHCP configuration

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **Import > Microsoft DHCP**. The **Import a Microsoft DHCP server dump** wizard opens.
4. Click on **[BROWSE]** to find the Microsoft dump file. Once you clicked on Open, the file is displayed in the field **File name**.
5. In the drop-down list **DHCP server**, select the target server.
6. In the section **Import global options**, tick the box if you want to apply options configured in the Microsoft DHCP dump to the destination server.
7. Click on **[OK]** to complete the operation. The report opens and closes.

Several Microsoft configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted. If two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

Note that, for large configurations, SOLIDserver runs the import process in the background. As it can take a while, the result is not displayed immediately.

Importing an Infoblox DHCP Configuration

You can import DHCP configurations from Infoblox solutions in IPv4. Through this import the whole DHCP server configuration creates the following elements within an EfficientIP DHCP server: scopes, ranges, leases, statics, groups and DHCP options. However, there are some restrictions:

- **Scopes restriction:** if the server you are importing contains overlapping scopes, only the first scope is imported, the rest is ignored.
- **Statics restriction:** statics associated to an IP address not included in one the scopes you are importing are ignored.
- **Shared network restriction:** shared network options are ignored.

- **DHCP options restriction:** only standard options are imported. If the server was configured using *non standard* DHCP options, they are imported only if they were previously defined either in the configuration file or within the SOLIDserver appliance.
- **Failover restriction:** Failover channels are not imported.
- **Infoblox options restriction:** all Infoblox options are ignored (these options usually include "infoblox" in their name).

The Infoblox DHCP loads its configuration from the file named *dhcpd.conf*. This file contains the whole configuration of the DHCP server. Within SOLIDserver:

- You must import your configuration at scope level within a specific DHCP server.
- You can import your configuration in a smart architecture as long as it is a One-to-One, a One-to-Many or a Single-Server smart architecture. It is impossible to import this configuration in a Split-Scope or a Stateless architecture.

To import an Infoblox DHCP configuration

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. In the menu, select **Import > Infoblox DHCP**. The **Import an Infoblox dhcpd.conf file** wizard opens.
4. Click on **BROWSE** to find the Infoblox dhcpd.conf file. Once you clicked on Open, the file is displayed in the field **File name**.
5. In the drop-down list **DHCP server**, select the target server.
6. Click on **OK** to complete the operation. The report opens and closes.

This procedure also works from the page All scopes of the server for which you want to import the Infoblox configuration: at server level click on the name of the server concerned, once on the page All scopes, follow the procedure from step 4. The server is automatically selected in the drop-down list DHCP server.

Several Infoblox configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted: if the configurations conflict with each other all the different elements are added. In the same way, if two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

Importing a MetaIP DHCP Configuration

You can import DHCP configurations from Meta Info DHCP software solution in IPv4. Through this import, the whole DHCP server configuration create the following elements within an EfficientIP DHCP server: scopes, ranges, leases, statics, groups and DHCP options. However, there are some restrictions:

- **Scopes restriction:** if the server you are importing contains overlapping scopes, only the first scope is imported, the rest is ignored.
- **Statics restriction:** statics associated to an IP address not included in one the scopes you are importing are ignored.
- **Shared network restriction:** shared network options are ignored.

- **DHCP options restriction:** only standard options are imported. If the server was configured using *non standard* DHCP options, they are imported only if they were previously defined either in the configuration file or within the SOLIDserver appliance.

The Meta IP DHCP loads its configuration from the file named *dhcpd.conf*. This file contains the whole configuration of the DHCP server. You can import this file directly from SOLIDserver GUI at scope level..

To import a Meta IP DHCP configuration

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **Import > Meta IP DHCP**. The **Import a Meta IP dhcpd.conf file** wizard opens.
4. Click on **[BROWSE]** to find the Meta IP dhcpd.conf file. Once you clicked on Open, the file is displayed in the field **File name**.
5. In the drop-down list **DHCP server**, select the target server.
6. Click on **[OK]** to complete the operation. The report opens and closes.

This procedure also works from the page All scopes of the server for which you want to import the Meta IP configuration: at server level click on the name of the server concerned, once on the page All scopes, follow the procedure from step 4. The server is automatically selected in the drop-down list DHCP server.

Several Meta IP configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted: if the configurations conflict with each other all the different elements are added. In the same way, if two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

Importing a Nortel NetID Configuration

You can import DHCP configurations from Nortel NetID solution in IPv4. SOLIDserver supports the following NetID configuration of DHCP:

- Scopes.
- Scope options.
- Ranges.
- Range options.
- Reservations.
- Reservations options.

The NetID DHCP loads its configuration from the file named *dhcpcfg.cur*. This file contains the whole configuration of the NetID DHCP server. SOLIDserver allows importing this file directly from its graphical user interface at scope level of the DHCP organization.

To import a NetID DHCP configuration

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.

2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **⬅ Import > NetID DHCP**. The **Import a NetID DHCP dump file** wizard opens.
4. Click on **[BROWSE]** to find the NetID dump file. Once you clicked on Open, the file is displayed in the field **File name**.
5. Click on **[NEXT]**. The page **Select DHCP server preferences** opens.
6. In the drop-down list **DHCP server**, select the target server.
7. Using the following drop-down lists, select the elements you want to retrieve from the configuration file:

Table 11.1. NetID DHCP import options

Option	Description
Scopes	Specify if you want to import the scopes from the configuration file.
Scopes options	Specify if you want to import the options tied to the scopes from the configuration file.
Ranges	Specify if you want to import the ranges from the configuration file.
Ranges options	Specify if you want to import the options tied to the ranges from the configuration file.
Statics	Specify if you want to import the reservations (statics) from the configuration file.
Statics options	Specify if you want to import the options tied to the reservations (statics) from the configuration file.

8. Click on **[OK]** to complete the operation. The report opens and closes.

Several NetID configuration files can be imported one after the other on the same target DHCP server. It allows to merge different DHCP configurations on one unique DHCP server. Through all the imports, no data is deleted: if the configurations conflict with each other all the different elements are added. In the same way, if two configuration files have a scope in common but named differently, the first scope name imported is overwritten by the new scope name.

Chapter 12. Importing DNS Data

SOLIDserver offers several ways of importing zones and RRs from legacy DNS servers to EfficientIP DNS servers. The DNS data can be downloaded or transferred from the GUI without having to install any tools on the remote system. You can import:

- [BIND archive files](#).
- [VitalQIP archive files](#).

Before importing, keep in mind that:

To import zones and records from a CSV File, refer to the section [Importing Data to the DNS](#).

Importing DNS Zones from a BIND Archive File

Within the DNS module you can import a BIND archive file containing all your zones.

Prerequisites

- The archive file **must be imported in a DNS server**, preferably an EfficientIP DNS server, an EfficientIP DNS Package server or a smart architecture.
- The archive file **must contain all the directories of your BIND configuration** including: the file *named.conf*, the zone files and any other necessary files whether they belong to the same directory or other sub directories.
- The archive file must have one of the following extensions: **.tar, .tgz, .gz or .zip** . It is not necessary to change the directory paths of your zone files in the file *named.conf* , if you are not able to provide the whole directory organizations in the archive file, the system can retrieve the files in the archive (several zone files may use the same name in different directories).

Limitations

- You **cannot use the characters "_", "@" and ":"** when importing a BIND archive file. Make sure you did not use any of these characters in zone names, record names, etc. as it would trigger either parsing errors (and not import the file) or import everything but the line containing these characters. For more details, refer to the RFC 1034 *Domain Names - Concepts and Facilities*.
- The archive must only contain supported BIND options. **Any non-supported BIND option declared in the archive file is ignored**. Once the archive file is imported, you can configure these extra options following the appendix [Configuring Non-Supported BIND Options](#).
- The file *named.conf* can contain include directives linking to other files if any **include directive is declared outside existing clauses**. Any include directive declared within existing clauses like *option {}*, *zone {}*, etc. is ignored. The file(s) declared in the directive include must be part of the archive file.
- The archive **cannot contain the directive \$GENERATE**, this directive is not supported.

Importing a BIND Archive File

Once you comply with the prerequisites and limitations, you can import your BIND archive file.

To import zones from a BIND archive file

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the name of the server in which you want to import the BIND archive file. The page **All zones** opens.
3. In the menu, select **Import > BIND archive file**. The wizard **Importing a BIND archive file** opens.
4. Click on **BROWSE** to select the BIND archive file to import.
5. In the field **File name**, the file is displayed.
6. In the drop-down list **DNS Server**, select the server that should receive the configuration. The server you click on is automatically selected.
7. In the drop-down list **Action**, you can either *Import data* or *Check file*.
8. You can tick the box **Import all options** if you want to import the global configuration settings that apply to all the zones.
9. Click on **OK** to complete the operation. The report opens and works for a while before displaying the import result and potential errors.
10. In the section **Export format**, you can download the import result report in **TEXT**, **HTML** or **EXCEL**.
11. Click on **CLOSE** to go back to the page **All servers**.

Importing the Content of a BIND Zone

If you already created your BIND zone in the GUI, you can import its content from the page **All RRs** of the zone.

This import relies directly on the BIND zone file itself, `<zone-name>`, without extension located in the DNS database.

To import the content of a BIND zone

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the name of the server in which you want to import the BIND archive file. The page **All zones** opens.
3. Click on the name of the zone. The page **All RRs** opens.
4. In the menu, select **Import > BIND zone file**. The wizard **Importing a BIND zone file** opens.
5. Click on **BROWSE** to select the BIND zone file to import.
6. In the field **File name**, the file is displayed.
7. Click on **OK** to complete the operation. The report opens and closes. The resource records are listed.

Importing DNS Zones from a VitalQIP Archive File

From the DNS page *All servers* you can import VitalQIP data in a `.qef` file. Keep in mind that this file includes and imports VitalQIP DNS as well as IPAM data at the same time.

To import VitalQIP data into SOLIDserver:

- The QIP export file (**.qef*) must belong to an archive file which extension is *.tar*, *.gz*, *.tgz*, *.tar.gz* or *.zip*.
- The *.qef* file must be located at the root of the archive file.
- The archive file does not need to include the files **_aud.qef* as they are not relevant to the import and might make your import take longer.

To import zones from a VitalQIP archive file

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **Import > QIP data**. The wizard **Importing a QIP file** opens.
3. Click on **BROWSE** to select the QIP archive file to import.
4. In the field **File name**, the file is displayed once selected.
5. Click on **NEXT**. The last page of the wizard opens.
6. In the drop-down list **DNS Server**, select the server that should receive the configuration.
7. Click on **OK** to complete the operation. The report opens and works for a while before displaying the import result and potential errors.
8. In the section **Export format**, you can download the import result report in **TEXT**, **HTML** or **EXCEL**.
9. Click on **CLOSE** to go back to the page **All servers**.

Chapter 13. Exporting Data

Within SOLIDserver, exporting data follows a set of rules:

- **The object parameters that you can export correspond to the columns of the page**

That way, on the one hand you can export the name of the object container: if you export a list of zones you can also export the name of the server and view they belong to. And on the other hand, you can export the customized parameters that you created through Class Studio and displayed as columns. These columns are preceded by the mention *Class param*: in the wizard.

- **An export is generated one level at a time**

If you are exporting zones from the page All zones in the DNS, you only export the zones themselves but not the RRs they contain.

- **An export can be generated in five different formats**

You can export lists of objects in .csv, .html, .xml, .xls and .pdf¹. Only the .csv file format provides the possibility to reimport the list again in the GUI.

- **An export can take into account from 1 to n objects**

On any page, exporting data takes into account every object listed. However, if you tick one or more elements, only the parameters of the ones you ticked are exported.

- **An export can be done at a specific time or scheduled to be generated regularly**

From the export wizard, you can choose to export the data right away or later on, even on a regular basis and at the frequency of your choosing.

- **An export name provides time and format information**

An export is always named after its format and moment of generation, never after what it contains. Each export is named as follows: *export_<extension>_<date>_<time>.<extension>*. Where *extension* refers to the export format; *date* is displayed as such: YYYYMMDD and time as such: HHMMSS. For instance, "export_excel_20130301_073042.xls" is an export generated in EXCEL on March 1st, 2013 at 07:30:42.

- **If the page does not have the menu Report, you cannot export the data listed**

Within SOLIDserver, almost any page allows to export data. To see the whole list of pages where you can export data, refer to the section [Pages where you can export data](#) below.

All exports are displayed on a single page, however the configuration files of the scheduled exports are displayed on their own page.

You can export data from almost any page. The menu *Report* indicates which pages are concerned:

Table 13.1. Pages where you can export data

Module	Pages
IPAM	All the pages of the module allow data exports
DHCP	All the pages of the module allow data exports

¹Keep in mind when exporting data to a PDF file that the number of columns is limited to 40 and affects the final display and might generate a file very hard to read.

Module	Pages
DNS	All the pages of the module allow data exports
NetChange	All network devices
	All VLANs
	All ports
	All routes
	All configurations
	All discovered items
Device Manager	All the pages of the module allow data exports
VLAN Manager	All the pages of the module allow data exports
Administration	Centralized Management
	Groups
	Users
	Syslog
	Session tracking
	User tracking
	Alerts
	Custom database
	Custom data

The Export Wizard

No matter what format you chose for the export, the **Export <format> file** wizard opens and looks like the image below.

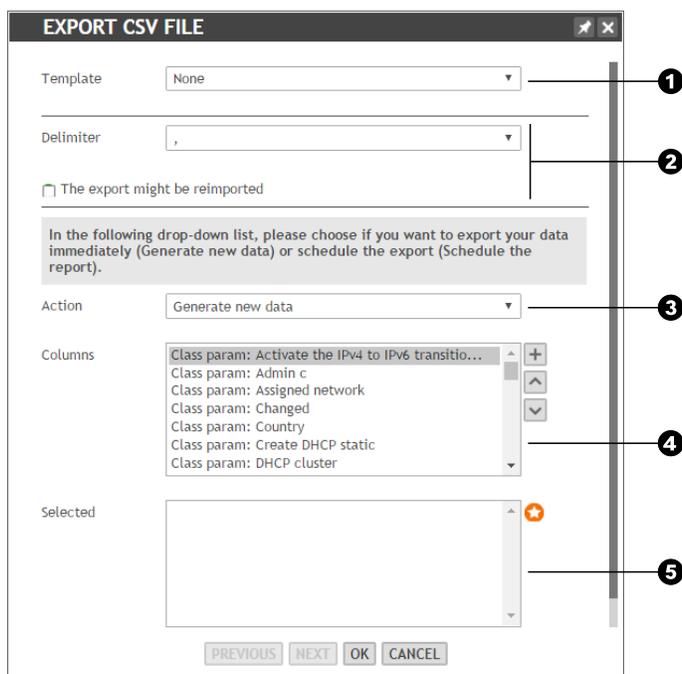


Figure 13.1. The CSV export wizard

- ❶ **Template** is a drop-down list that allows you to save all your configuration as a template for later exports of the list.
- ❷ When exporting CSV files, you can find two extra fields. First, the drop-down list **Delimiter** allows you to select which delimiter you want to use during the data export. Second, the box **The export might be reimported**² can be ticked if you want to reimport the data in a SOLIDserver appliance: this basically exports the list as raw data that is easier (and therefore faster) to reimport.
- ❸ **Action** is a drop-down list that allows you to export right away your list or schedule the export it at the frequency of your choice.
- ❹ **Columns** is a list that allows you to select the columns, i.e. parameters, of your choice. This list contains all the columns that you can display on the page as well as the class parameters related to the objects of the list.
- ❺ **Selected** is a list that sums up all the columns that you selected and which data you are about to export. It also allows you to order the data according to your needs.

Browsing the Exports Database

All the exports must be downloaded at the end of the export wizard.

Only the scheduled exports are available on the page *Local files listing*. The schedule configuration is available on the page *Scheduled exports*.

To display the scheduled exports

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens. By default, it displays the list *Local* where you can find all your exports.

To display the scheduled exports configuration

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Exports*, click on **Scheduled**. The page **Scheduled exports** opens.

Configuring Exports

The export can be of numerous forms as you can choose an export format, to schedule it or not and finally save your columns configuration in a template and later on use the template as is or use it as a basis during another export.

To export data immediately

1. Go to the page of your choice.
2. Tick the objects of your choice or none if you want to export the whole list.
3. In the menu, select  **Report** > **Export** > *<format of your choice>*. The wizard **Export <format> file** opens.
4. In the drop-down list **Template**:
 - a. Select *None* if you do not want to create a template and go to step 5.

²Exporting data without ticking this box might trigger some errors. Some columns might not be imported at all, for instance the value of the column *Network is terminal* cannot be imported if the box is not ticked.

- b. Select *New template* if you want to create a template. The field **Template name** appears, name your template. The template saves the columns you select as well as the delimiter if you export the list in a .csv file.

Once the export is generated, you can rename or delete the templates if need be. For more details, refer to the section [Managing Export Templates](#) below.

5. If you chose to export a CSV file:
 - a. In the drop-down list **Delimiter**, select *comma*, *semi-colon* or *tab*.
 - b. Tick the box **The export might be reimported**³ to export the list or selected objects as raw data.
6. In the drop-down list **Action**, select *Generate new data*.
7. In the list **Columns**, select one by one the columns that you want to export and click on . They are moved to the list **Selected**.
8. In the list **Selected**, you can order the columns according to your needs using  and . To remove a column from the export, select it and click on . It is moved back to the list **Columns**.
9. Click on to complete the operation. The report opens and works for a while.
10. You can click on to save the export. The page refreshes when the export is over.
11. Click on to close the wizard. The page is visible again.

From the menu *Report* you can also schedule exports. Keep in mind that these exports are managed differently: the generated file is available in the *Local files Listing*, plus scheduling an export creates a configuration that you can manage on the page *Scheduled exports*. For more details, refer to the section [Managing Scheduled Exports Configuration Files](#) below.

To schedule an export

1. Go to the page of your choice.
2. Tick the objects of your choice or none if you want to export the whole list.
3. In the menu, select  **Report** > **Export** > *<format of your choice>*. The wizard **Export <format> file** opens.
4. In the drop-down list **Template**, you can:
 - a. Choose not to create a template by selecting *None* and export your data.
 - b. Choose to create a template by selecting *New template*. The field **Template name** appears, name your template. The template saves the columns you select as well as the delimiter if you export the list in a .csv file.
5. If you chose to export a CSV file:
 - a. In the drop-down list **Delimiter**, select *comma*, *semi-colon* or *tab*.
 - b. In the section **The export might be reimported**, check the box to export the list or selected objects as raw data.
6. In the drop-down list **Action**, select *Schedule the report*. The page refreshes.

³This option must be ticked if you plan on reimporting some data, for instance the value of the Terminal column.

7. In the list **Columns**, select one by one the columns that you want to export and click on . They are moved to the list **Selected**.
8. In the list **Selected**, you can order the columns according to your needs using and . To remove a column from the export, select it and click on . It is moved back to the list **Columns**.
9. Click on . The last page of the wizard opens.
10. Configure the export frequency or date and time (UTC) of the export using the table below.

Table 13.2. Scheduled export fields

Field	Description
Day(s) of the week	In this drop-down list, select a frequency (over the whole week or for a specific set of days) or a specific day of the week. By default, <i>Every day</i> is selected.
Date of the month	In this drop-down list, select a specific day of the month or a frequency (every day) for the refresh. By default, <i>Every day</i> is selected.
Month	In this drop-down list, select a specific month or a frequency (every month) for the refresh. By default, <i>Every month</i> is selected.
Hour	In this drop-down list, select a frequency (over the whole day or for a limited period of time each day), a set of hours or a specific hour per day for the refresh. The hour respects the UTC standard. By default, <i>Every hour</i> is selected.
Minute	In this drop-down list, select the moment (o'clock, quarter past, half past or quarter to) or the frequency (in minutes) of the refresh. The minute respects the UTC standard. By default, <i>Every minute</i> is selected.
Name	In this field, name the scheduled export in this field.
Mail to	In this drop-down list, select the group which users should receive the export notification email. This email cannot be sent if the users email address is not valid or if your SMTP relay is not configured. For more details, refer to the section Configuring the SMTP Relay . By default, the first of your groups, in the ASCII alphabetic order, is selected.
Rights as	In this drop-down list, select a user. His/her rights and limitations are applied in the report: only the items this user has access to are listed in the export.

11. Click on to complete the operation. The report works and displays the export report.
12. Click on to go back to the page.

The export configuration is available on the page *Scheduled exports*. For more details, refer to the procedure [To display the scheduled exports configuration](#).

When the export is generated, it is available on the page *Local files listing*. For more details, refer to the procedure [To display the scheduled exports](#).

Exporting Data To Reimport It Later

Any list you might have exported can be reimported on most pages as long as you exported it in a CSV file. For more details, refer to the table [IPAM pages where you can import CSV files](#). **We strongly recommend that during the export you tick the box** *The export might be reimported*, to make it faster to process.

Keep in mind that in each of these modules and pages, you can reimport all the parameters of your choice but some columns are required, and without them, you cannot go through with the import. So when exporting, you must select these columns. In the sections below, we will only detail the pages where you can actually import data.

Required Columns To Reimport Data in the IPAM Module

The export wizard is accessible in the menu *Report* of any page of the module IPAM. On each page, some data is required during an import so you might need the columns listed in the table below if you intend to reimport a CSV file.

Keep in mind that the field **Space name** of the import wizard allows you to:

- Select the corresponding column of your CSV file
- Select one space among the ones in your database or the option *Use best space*, with IPv4, the option uses the IP address and size to place the object in the best space, block-type network and/or subnet-type network possible.

Table 13.3. Required columns to reimport data in the IPAM

Object	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
On the page All spaces			
Space	Name	Space Name	Name
On the page All networks			
Block-type network	Start	Network start address	First address
	Space	Space name	Space name
Block-type network (v6)	Start	Network start address	First address
	Prefix	Network prefix	Prefix
	Space	Space name	Space name
Subnet-type network	Address	Address + prefix ^a	Address
	Name	Network name	Name
	Space	Space name	Space name
Subnet-type network (v6)	Address	Network address	Address
	Prefix	Network prefix	Prefix
	Name	Network name	Name
	Space	Space name	Space name
On the page All pools			
Pool	Start address	Pool start address	First address
	Name	Pool name	Name
	Space	Space name	Space name
Pools (v6)	Start address	Pool start address	First address
	End address	Pool end address	Last address
	Name	Pool name	Name
	Space	Space name	Space name
On the page All addresses			
IP address	Address	IP address	IP address
	Name	IP name	Name
	Space	Space name	Space name
IP address (v6)	Address	IP address	IP address
	Name	IP name	Name
	Space	Space name	Space name

^aThis field can be used to export and reimport the network start address and size.

Required Columns To Reimport Data in the DHCP Module

The export wizard is accessible in the menu *Report* of any page of the module DHCP, except the pages *All servers*, *All leases* and *All leases (v6)*. On each page, some data is required during an import so you need the columns listed in the table below if you intend to reimport a CSV file.

Keep in mind that the **DHCP server** and **DHCP6 server** fields of the import wizard allow you to:

- Select the corresponding column of your CSV file
- Select one server among the ones in your database.

Table 13.4. Required columns to reimport data in the DHCP

DHCP page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All scopes	Address	DHCP scope address	Address
	Server	DHCP server name	DHCP server
All ranges	Start address	DHCP range start addr	Start address
	Server	DHCP server name	DHCP server
All statics	Name	DHCP static name	DHCP static name
	MAC address	MAC address	MAC address
	Server	DHCP server name	DHCP server
All scopes (v6)	Address	Address	Start address
	Server	Server	DHCP6 server
All ranges (v6)	Start address	Start address	Start address
	End address	End address	End address
	Server	Server	DHCP6 server
All statics (v6)	Name	DHCP static name	DHCP static name
	Server	Server	DHCP6 server

Required Columns To Reimport Data in the DNS Module

The export wizard is accessible in the menu *Report* of any page of the module DNS. On each page some data, is required during an import so you need the columns listed in the table below if you intend to reimport a CSV file.

Table 13.5. Required columns to reimport data in the DNS

DNS page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All zones	Name	Zone name	DNS zone name
	Type	Zone type	DNS zone type
	Server	DNS server name	DNS server name
All RRs	RR name	RR name	RR name
	Value	RR value	Value 1 ^a
	Zone	DNS zone name	Zone name
	Type	Space name	RR type

^aThis field includes all the information exported from the column *Value*.

Required Columns To Reimport Data in NetChange Module

The export wizard is accessible in the menu *Report* of any page of the module NetChange. On each page, some data is required during an import so you need the columns listed in the table below if you intend to reimport a CSV file.

Table 13.6. Required columns to reimport data in NetChange

NetChange page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All network devices	Space	Space Name	Target space

Required Columns To Reimport Data in Device Manager Module

The Export wizard is accessible in the menu *Report* on of any page of the module Device Manager. On each page, some data is required during an import so you need the columns listed in the table below if you intend to reimport a CSV file.

Table 13.7. Required columns to reimport data in Device Manager

Device Manager page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All devices	Name	Device name	Name
All ports & interfaces	Name	Interface name	Name
	Type	Type	Type
	Space	Device name	Device

Required Columns To Reimport Data in VLAN Manager Module

The Export wizard is accessible in the menu *Report* on on any page of the VLAN Manager module. Some data is required on each page during an import, therefore, you might need to specify the columns listed in the table below if you intend to reimport a CSV file.

Table 13.8. Required columns to reimport data in VLAN Manager

VLAN Manager page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All domains	Name	Name	Name
	Start ID	Domain Start ID	Start ID
	End ID	Domain End ID	End ID
All ranges	Name	Name	Name
	Start ID	Range Start ID	Start ID
	End ID	Range End ID	End ID
	Domain	Range Domain	Domain
All VLANs	VLAN ID	VLAN ID	VLAN ID
	Range	Range	Range
	Domain	Domain	Domain

Required Columns To Reimport Data in VRF Module

The Export wizard is accessible in the menu on of any page of the module VRF. On each page, some data is required during an import so you might need the columns listed in the table below if you intend to reimport a CSV file.

Table 13.9. Required columns to reimport data in VRF

VRF page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
All VRFs	Name	VRF name	VRF name
	RD	VRF RD	VRF RD
All VRF Route Targets	Source RD	Source RD	Source RD of the VRF Route Targets
	Target RD	Target RD	Target RD of the VRF Route Targets

Required Columns To Reimport Data in the Administration Module

The Export wizard is accessible in the menu *Report* on a limited number of pages of the Administration module: Centralized Management, Groups, Users, Session tracking, User tracking, Alerts, Custom database and Custom data.

You can import, or reimport, data on three of these pages.

Table 13.10. Required columns to reimport data on the module Administration

Administration page	Listing page required column(s)	Column name in the export wizard	Column name in the import wizard
Groups	Name	Name	Name
Users	Login	Login	Login
Custom data	<i>First column</i>	/	Value 1

Managing Scheduled Exports

Once generated, all the scheduled exports are saved in the directory `/data1/exports`. In the GUI, they are available in the page **Local files listing**, under the filter page **Local**, where you can export or delete them.

Each column on the page corresponds to the parameters configured during the export configuration. You can sort the list through each column, you can filter it through the columns *Name*, *Type* and *Owner*. You cannot edit the listing layout of this page or access a properties page as all the information is displayed.

To download a scheduled export from the Local files listing

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens. By default, it displays the list *Local* where you can find all your scheduled exports.
3. Click on the name of the export of your choice to download it.

To delete a scheduled export from the Local files listing

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens. By default, it displays the list *Local* where you can find all your scheduled exports.
3. Tick the export(s) you want to delete.
4. In the menu, select  **Edit** > **Delete file(s)**. The wizard **Delete file** opens.

5. Click on **OK** to complete the operation. The report opens and closes. The page refreshes, the selected scheduled export is no longer listed.

Managing Scheduled Exports Configuration Files

If you created scheduled exports, the configuration file is on the page *Scheduled exports*. Once created a scheduled export, you cannot edit its configuration. However, you can disable and enable it or even delete it.

All the configuration files are listed and each column corresponds to the parameters configured during the scheduled export creation. You can sort and filter the list through each column but you cannot edit the listing layout of this page. The scheduled exports do not have a properties page as all the information is displayed.

To enable/disable a scheduled export configuration file

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Exports*, click on **Scheduled**. The page **Scheduled exports** opens.
3. Tick the configuration file(s) of your choice.
4. In the menu, select **Edit > Enable** or **Disable**. The wizard opens.
5. Click on **OK** to complete the operation. The report opens and closes. The file is marked **OK** or **Disabled**.

To delete a scheduled export configuration file

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Exports*, click on **Scheduled**. The page **Scheduled exports** opens.
3. Tick the configuration file(s) you want to delete.
4. In the menu, click on **Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The page refreshes, the file is no longer listed.

Managing Export Templates

Every export template you created is listed on the page **Import/Export Templates Management** of each module. On this page, you can either rename or delete the templates.

This page is composed of as many panels as there are pages where you can import data in the module. Each panel lists all the templates configured on the page, the export templates are listed as follows: *Export: <template_name>*.

To rename an export template

1. In the sidebar, go to the module of your choice.
2. In the menu, select **Extra options > Import/Export templates management**. The page **Import/Export Templates Management** opens.
3. In the panel of your choice, select the *Export: <template_name>* you want to rename.

4. Click on **RENAME**. The wizard **Rename template** opens.
5. In the field **New Name**, rename your template.
6. Click on **OK** to complete the operation. The report opens and closes. The name changes in the list.

To delete an export template

1. In the sidebar, go to the module of your choice.
2. In the menu, select **⋮** **Extra options** > **Import/Export templates management**. The page **Import/Export Templates Management** opens.
3. In the panel of your choice, select the *Export: <template_name>* you want to delete.
4. Click on **DELETE**. The wizard **Delete template** opens.
5. In the field **New Name**, rename your template.
6. Click on **OK** to complete the operation. The report opens and closes. The template is no longer listed.

Part IV. Dashboards

Dashboards is the first module you see when you connect to SOLIDserver.

The *Main Dashboard* is the appliance homepage. For the superuser, it provides an overview of the appliance configurations and services.

The dashboards are organized per module. On each one you can gather gadgets to monitor data or set up custom shortcuts and search engines to ease up the management.

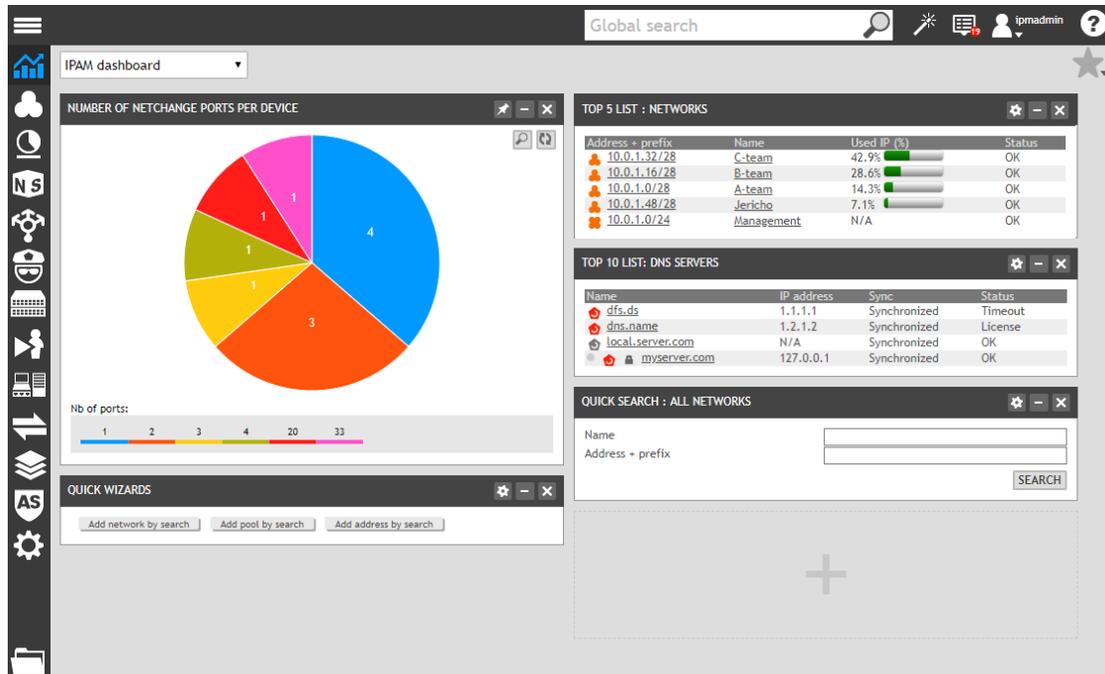


Figure 71. The IPAM dashboard

This part contains the following chapters:

- [Building Dashboards](#) describes all the customization options available from the dashboards.
- [Managing Gadgets](#) describes all the gadgets that you can create or display on the dashboards.

Chapter 14. Building Dashboards

Most modules have a dedicated dashboard where you can add and organize gadgets to set up a customized display of information. For more details regarding gadgets, refer to the chapter [Managing Gadgets](#).

Browsing the Dashboards

All dashboards are accessible from a single drop-down list.

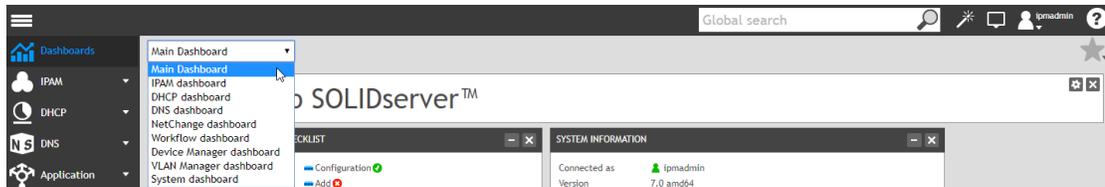


Figure 14.1. The dashboard selection drop-down list

To display a dashboard

1. In the sidebar, click on **Dashboards**. The page refreshes.
2. In the drop-down list, select one of the available dashboards detailed below. The page refreshes.

Table 14.1. Available dashboards

Dashboard	Description
Main Dashboard	SOLIDserver homepage. By default it displays a set of gadgets that may differ depending on the user connected. For more details, refer to the section Gadgets Displayed by Default .
IPAM dashboard	The dashboard of the module IPAM. By default, it is empty.
DHCP dashboard	The dashboard of the module DHCP. By default, it is empty.
DNS dashboard	The dashboard of the module DNS. By default, it is empty.
NetChange dashboard	The dashboard of the module NetChange. By default, it displays a set of gadgets containing network device related charts. For more details, refer to the section Gadgets Displayed by Default .
Workflow dashboard	The dashboard of the module Workflow. By default, it is empty.
Device Manager dashboard	The dashboard of the module Device Manager. By default, it displays a set of gadgets containing device related charts and a Top list. For more details, refer to the section Gadgets Displayed by Default .
VLAN Manager dashboard	The dashboard of the module VLAN Manager. By default, it is empty.
System dashboard	The dashboard of the module Administration. By default, it is empty.

Note that the modules Application, Guardian, VRF and SPX do not provide a dashboard.

Assigning a Gadget from the Dashboard

From any dashboard you can assign existing gadgets using a button located at the bottom of the dashboard.

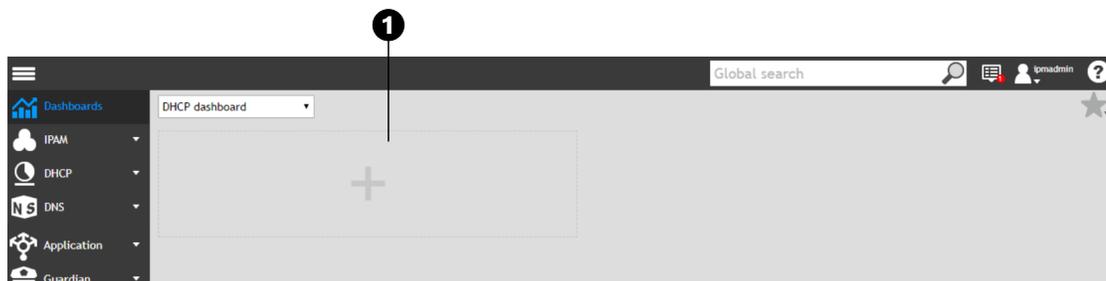


Figure 14.2. The button Add a Gadget available on any dashboard

- 1 The button **+** opens the wizard **Add a gadget** where you can select an existing gadget to display on the dashboard.

To assign a gadget from a dashboard

1. Go to the dashboard of your choice.
2. If many gadgets are displayed, you may need to collapse **▣** some.
3. Click on **+**. The wizard **Add a gadget** opens.
4. In the list **Type**, select *Chart*, *Top List*, *Quick Search* or *Other*, that contains the default descriptive and configuration gadgets, and the gadget *Bookmarks* and *Quick Wizards* if you created them. For more details, refer to the section [Creating Gadgets](#).
5. Click on **NEXT**. The list **Gadget** displays the available gadgets of the selected type.
6. Select the gadget you want. If there is no gadget of this type yet, the list is empty.
7. Click on **OK** to complete the operation. The gadget is now visible on the dashboard.

If you no longer want to display a gadget on a dashboard, refer to the section [Hiding Gadgets from a Dashboard](#).

If you want to create, edit, disable, control the visibility or delete a gadget, refer to the chapter [Managing Gadgets](#).

Organizing Gadgets on a Dashboard

On any dashboard, you can organize the gadgets to suit your needs.

Moving a Gadget

You can **drag and drop** the gadgets to place them elsewhere on each dashboard.

When you put the mouse pointer in the gadget drag bar, it changes shape and you can see the former position of the gadget and how much space it takes up in the new spot.

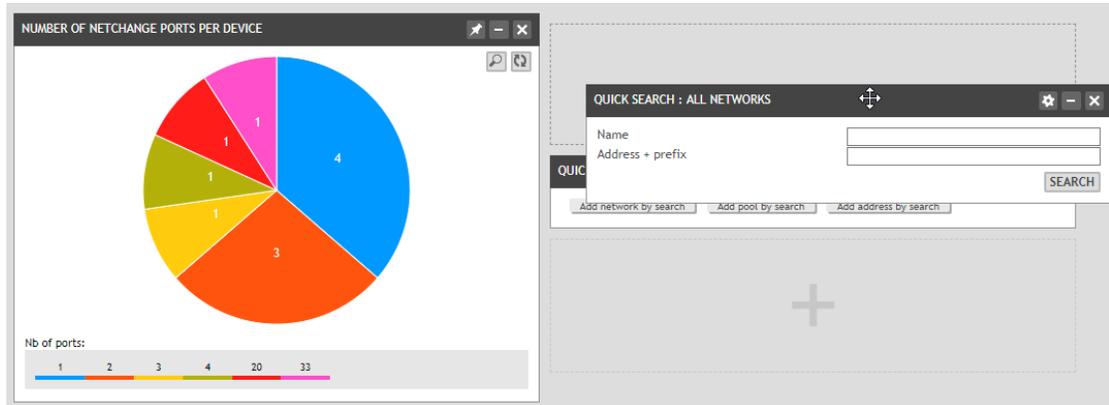


Figure 14.3. Moving a gadget

Collapsing or Expanding a Gadget

Next to the gadget name, you can use the buttons to **collapse**  a gadget.



Figure 14.4. Collapsing a gadget

Once collapsed, the gadget is displayed as a simple line. Only its drag bar is visible, it contains its name and buttons.

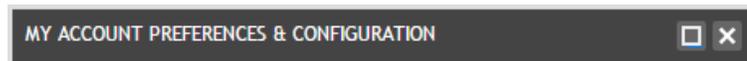


Figure 14.5. An example of a collapsed gadget

Next to the gadget name, you can use the buttons to **expand**  the gadget again.

Hiding Gadgets from a Dashboard

You can hide the gadgets displayed on a dashboard at any point.

Hiding a gadget means that it is no longer visible on the dashboard, it does delete it. To display it again, refer to the section [Assigning a Gadget from the Dashboard](#).

To hide a gadget from a dashboard

1. Go to the dashboard of your choice.
2. In the gadget drag bar, where the gadget name is displayed, click on . The wizard Disable gadget opens.
3. Click on to complete the operation. The wizard closes. The gadget is no longer visible.

Chapter 15. Managing Gadgets

The gadgets allow you to monitor the appliance and customize your dashboards. You can create, assign, hide and/or delete gadgets. Some can even be edited.

All the gadgets are composed of two parts:

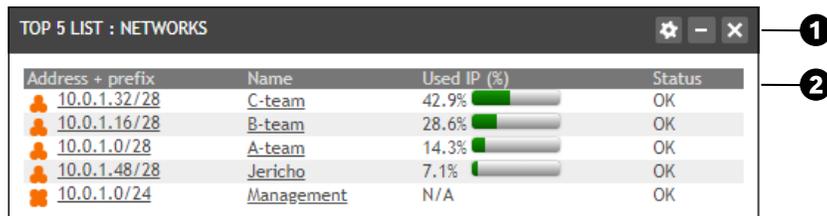


Figure 15.1. The common structure of the gadgets

- 1 The upper gray part is the gadget **drag bar**. It contains the gadget name and the buttons to collapse, expand or hide it. On some gadgets, the button allows to edit them.
- 2 The lower white part contains the information. Its content differs for every type of gadget.

You can manage existing charts like gadgets. By default, a set of gadgets are available on the page the *Gadgets Library*, for more details, refer to the appendix [Default Gadgets](#) and .

Browsing Gadgets

The gadgets are available on three pages:

- **My Gadgets** where you can manage all the gadgets already assigned to at least one dashboard.
- **Gadgets Library** where you can manage all the existing gadgets.
- **System statistics** that contains the appliance statistics. Every chart on the page can be used as a gadget and assigned to a dashboard.
- **Any properties page containing charts**. For more details, refer to the section [Assigning a Gadget from a Resource Properties Page](#).

A number of gadgets are displayed by default on some dashboards, as detailed in the section [Gadgets Displayed by Default](#).

Browsing the Assigned Gadgets

From the page *My Gadgets*, you can hide, display and manage the visibility of gadgets already assigned to a dashboard.

To display the page *My Gadgets*

1. From any page, in the top bar, select **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. To display the gadgets of a specific dashboard, in the search engine of the column **Dashboard**, type in the name of the dashboard of your choice.

The page contains the following columns, you cannot edit the list layout.

Table 15.1. The columns of the pages Gadgets Library and My Gadgets

Column	Description
Name	The gadget name.
All users	The gadget visibility: <i>Yes</i> means it is visible to all users, <i>No</i> means it is only visible for the user who created it.
Type	The gadget type: <i>Chart</i> , <i>Configuration</i> , <i>Descriptive</i> , <i>Quick Search</i> , <i>Shortcut</i> or <i>Top List</i> . For more details, refer to the section Understanding the Gadget Types .
Dashboard	The name of the dashboard(s) where the gadget is assigned.
Status	The gadget status: <i>Enabled</i> or <i>Disabled</i> .

Browsing the Gadgets Database

From the page *Gadgets Library* you can enable, disable and delete all the existing gadgets, whether they are already assigned to a dashboard or not.

It contains the same columns as the page *My Gadgets*.

To display the page Gadgets Library

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.

To display a gadget assignation details from the page Gadgets Library

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. In the column **Name**, click on the gadget of your choice. The page **My Gadgets** opens and lists the gadget as many times as it has been assigned. In the column *Dashboard*, the name of the dashboard where the gadget is assigned is listed once. If the list is empty, it means that the gadget is not assigned.

To display a gadget properties page from the page Gadgets Library

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. At the end of the line of the gadget of your choice, click on . The properties page opens.

The panel **Main properties** sums up the gadget name, type, visibility and status. If it contains the button [EDIT](#), you can edit the gadget. For more details, refer to the section [Editing Gadgets](#).

Browsing the System Statistics

Within the module Administration, the page *System statistics* provides panels containing charts that can be assigned to any dashboard as gadgets.

These charts are not listed on the page *Gadgets Library*, but once assigned they are listed on the page *My Gadgets*.

To display the local SOLIDserver Statistics

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **System statistics**. The page **System statistics** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, make sure your local appliance¹ is selected.

All the charts on the page, except *Processes state*, can be used as a gadget. The charts return:

- Traffic information, in the panels *DNS traffic*, *DHCP traffic*, *HTTP traffic*, *SNMP traffic* and *Database replication traffic*.
- System information, in the panels *Load average*, *CPU per process*, *Memory usage per process*, *Disk operations*, *I/Os per process*, *SQL queries*, *Threads*, *User sessions* and *Disk Usage*.

To display a chart on a dashboard, refer to the section [Assigning a Gadget from the Page System Statistics](#).

Keep in mind that these charts are empty during the first appliance use, without any traffic, there is no data to display.

Understanding the Gadget Types

There are four types of gadgets that you can create:

- [Chart](#), used to get a graphical overview of the data.
- [Top List](#), used to list resources based on specific filters.
- [Quick Search](#), used to create shortcuts for columns search engines.
- [Shortcut](#), used to create shortcut links toward specific pages. They include:
 - [The Gadget Quick Wizards](#): used to create shortcuts toward your quick wizards.
 - [The Gadget Bookmarks](#): used to create shortcuts toward your bookmarked pages.

In addition, there are two gadget types, the *Descriptive* and *Configuration* gadgets. You cannot create gadgets of these types, they are displayed by default. For more details, refer to the section [Gadgets Displayed by Default](#).

¹You can display the statistics of a remote appliance on the page but these charts cannot be used as gadgets.

Chart

The chart allows to compare values and labels and create graphical representations of resources activity or repartition via a pie chart or a bar chart. For instance, the DNS resource records type distribution among the zones of a server.

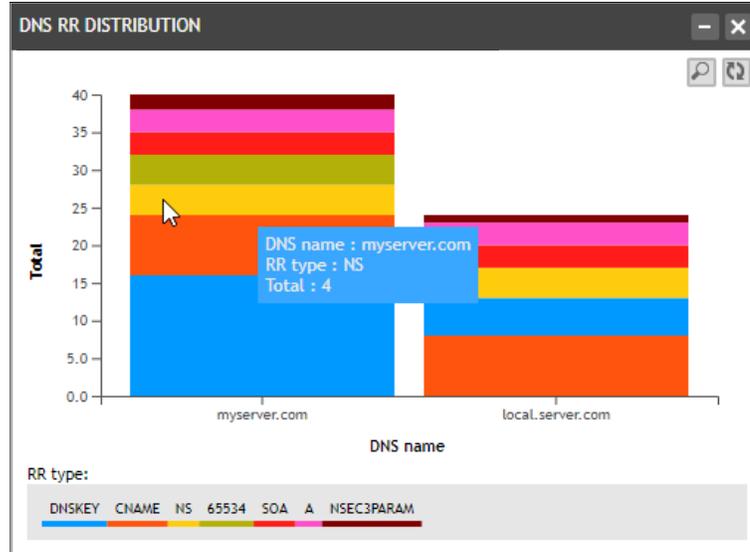


Figure 15.2. Example of a bar chart representation of a record types distribution

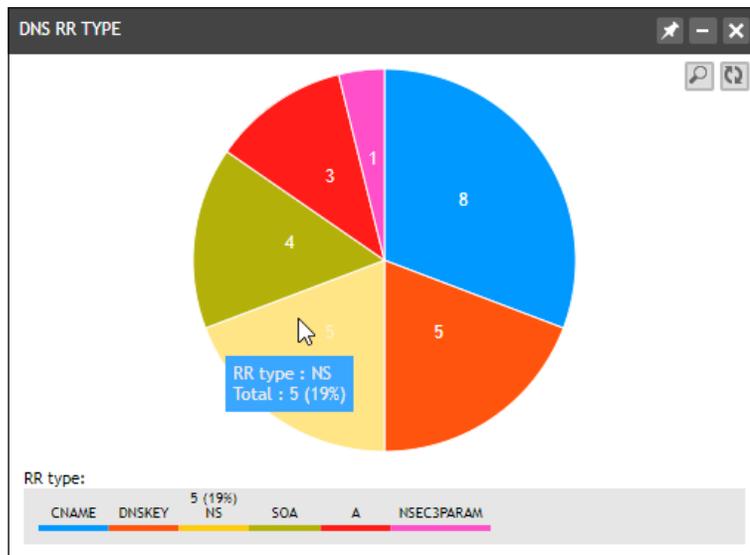
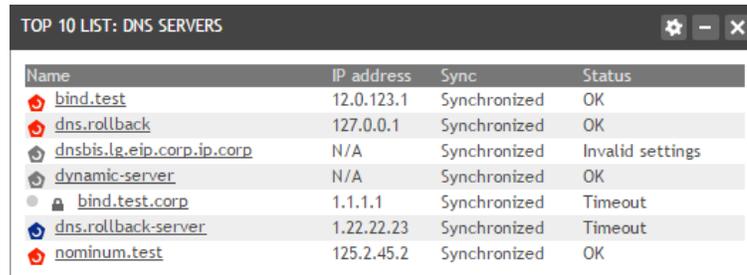


Figure 15.3. Example of a pie chart representation of a record types distribution

For more details, refer to the section [Creating a Chart Gadget](#).

Top List

The Top List allows you to create a specific list displaying from 5 to 25 items from a listing page. Once created, it looks like a table composed of a maximum of 4 columns that display the entries you want. They can match a filtered display of the source page. For instance, a list of the five most heavily used terminal networks.



Name	IP address	Sync	Status
bind.test	12.0.123.1	Synchronized	OK
dns.rollback	127.0.0.1	Synchronized	OK
dnsbis.lg.eip.corp.ip.corp	N/A	Synchronized	Invalid settings
dynamic-server	N/A	Synchronized	OK
bind.test.corp	1.1.1.1	Synchronized	Timeout
dns.rollback-server	1.22.22.23	Synchronized	Timeout
nominum.test	125.2.45.2	Synchronized	OK

Figure 15.4. An example of a Top List

The Top List can be edited thanks to the button . For more details, refer to the sections [Creating a Top List Gadget](#) and [Editing Gadgets](#).

Quick Search

The Quick Search allows you to create a filtering tool based on the columns of a page. It is composed of fields matching the columns search engines in the target listing page.

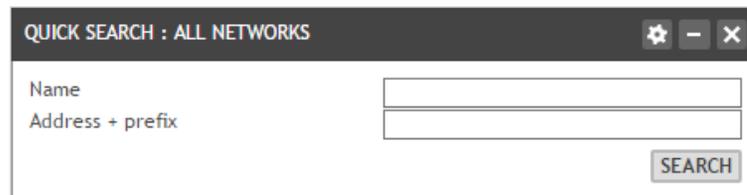


Figure 15.5. An example of a Quick Search

In the gadget, the selected columns are displayed as input fields. When you click on **SEARCH**, you execute a search that automatically opens the target page and applies the filters of the gadget to only return the matching results.

Thanks to this gadget, you can filter pages using several columns from any dashboard. For instance, you can search for specific zones through their name, type and DNSSEC configuration within a server from any module dashboard.

The Quick Search can be edited thanks to the button . For more details, refer to the sections [Creating a Quick Search Gadget](#) and [Editing Gadgets](#).

Shortcut

The Shortcut contains link buttons toward specific pages. They are only used for quick wizards and bookmarks. For more details, refer to the section [Shortcuts](#).

The Gadget Quick Wizards

This gadget contains links toward your quick wizards. It can only be displayed once on each dashboard.

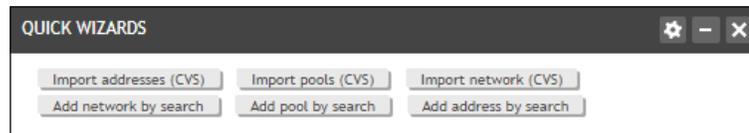


Figure 15.6. An example of a Quick Wizards gadget

You can customize the gadget content on each dashboard with . For more details, refer to the sections [Creating a Quick Wizards Gadget](#) and [Editing Gadgets](#).

The Gadget Bookmarks

This gadget contains links toward the bookmarks of your choice. You can display it on any dashboard.

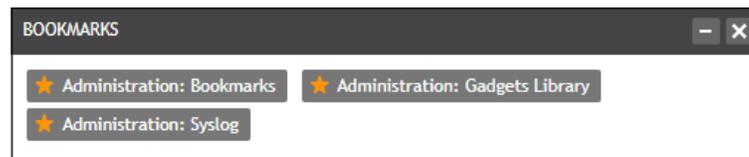


Figure 15.7. Example of a gadget Bookmarks

For more details, refer to the sections [Creating a Gadget Bookmarks](#) and [Editing Gadgets](#).

Gadgets Displayed by Default

The **superuser session**, the user *ipmadmin*, has a set of gadgets displayed by default:

- On the *Main Dashboard*: [System Information](#), [General Information](#), [SOLIDserver Configuration Checklist](#), [My Account Preferences & Configuration](#) and [Shortcuts](#).
- On *NetChange dashboard*: [NetChange Network Devices Vendor](#), [NetChange Active Ports Speed \(bps\)](#), [NetChange Port Status](#) and [Number of NetChange Ports per Device](#).
- On *Device Manager dashboard*: [Number of Ports Used per Device](#), [Number of Interfaces Used per Device](#) and [Alert on Ports/Interfaces Reconciliation Drift](#).

Keep in mind that on the session of **any other user**:

- Only [My Account Preferences & Configuration](#) is displayed on the *Main Dashboard*.
- Any default gadget can be displayed on the dashboards as detailed in the section [Assigning Gadgets from the Page Gadgets Library](#). They are all listed in the appendix [Default Gadgets](#).
- By default, the gadget [SOLIDserver Configuration Checklist](#) is only available for *ipmadmin*.

System Information

This descriptive gadget is available by default on the *Main Dashboard* of *ipmadmin* (the superuser session).

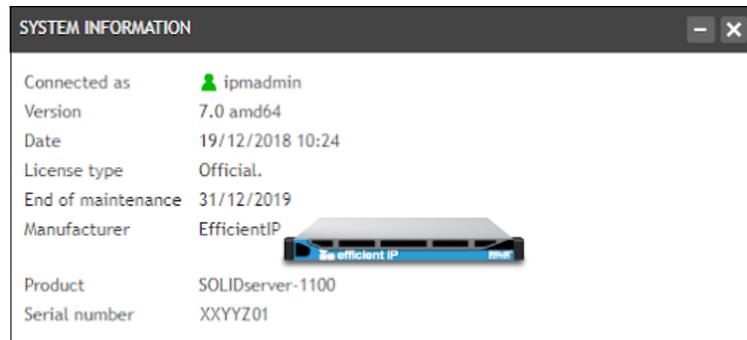


Figure 15.8. The gadget System information

This gadget sums up system and user related information:

Connected as

The name of the user connected. Click on the user name (*ipmadmin* in the image above) to open the wizard **Configure user settings** that allows to set the user account preferences. For more details, refer to the section [Account Configuration](#).

Version

The current software and architecture versions of SOLIDserver.

Date

The appliance current date and time. If it is not accurate, the *License type* and finally the support used (*Manufacturer*, *Product* and *Serial*).

License type

The appliance license type: *Temporary* (with the *End date* between brackets) or *Official* i.e. with no end date.

End of Maintenance

The date of the end of the appliance maintenance period.

Manufacturer

The appliance manufacturer name. It indicates if the appliance is installed on a virtual machine or a physical hardware appliance.

Product

The product name, either hardware (with its size) or software. It depends on the manufacturer.

Serial number

The appliance serial number. For hardware appliances, it is composed of 6 hexadecimal digits. For virtual appliances, it is only visible on the page *Centralized Management*.

General Information

This descriptive gadget is available by default on the *Main Dashboard* of *ipmadmin*.

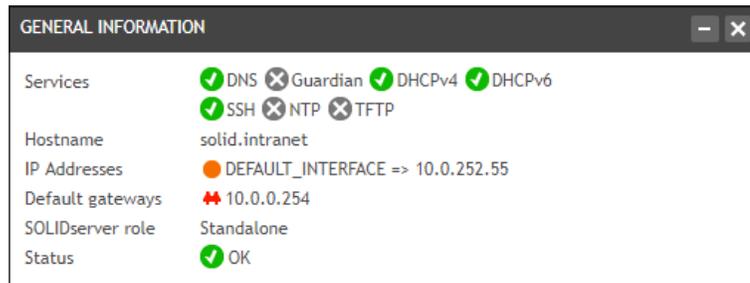


Figure 15.9. The gadget General information

This gadget sums up the current network and services configuration:

Services

The main services statuses: running ✔, disabled ✖ or not yet configured ✕. Click on each service name to manage them. A running service can be disabled from the dashboard. The services that are disabled or not configured yet provide a link to the page *Services Configuration*. For more details, refer to the chapter [Configuring the Services](#).

Hostname

The appliance hostname. Click on its name (*solid.intranet* in the image above) to open the page *Network Configuration* and edit it. For more details, refer to the chapter [Configuring the Network](#).

IP Addresses

The appliance interface(s) IP address. The default interface indicates the IP address you configured to connect to SOLIDserver. Click on the interface name (*DEFAULT_INTERFACE* in the image above) to open the page *Network Configuration* and edit it. For more details, refer to the chapter [Configuring the Network](#).

Default gateways

The appliance default gateway. Click on its IP address to open the page *Network Configuration* and edit it. For more details, refer to the chapter [Configuring the Network](#).

SOLIDserver role

The appliance role: *Standalone*, *Master* or *Hot Standby*. The last two roles imply that your SOLIDserver is configured in high availability (HA). Click on the role to open the page *Centralized Management*. For more details, refer to the chapter [Centralized Management](#).

Status

The appliance status. A *Standalone* appliance is always ✔ OK. The appliances configured in HA can be marked ✖, to indicate that the configuration is not working properly. Click on the status to open the page *Centralized Management*. For more details, refer to the chapter [Centralized Management](#).

SOLIDserver Configuration Checklist

This configuration gadget is by default displayed on the *Main Dashboard* of *ipmadmin*.

Only *ipmadmin* can share it to allow other users to display it, even if they belong to the group *admin*.

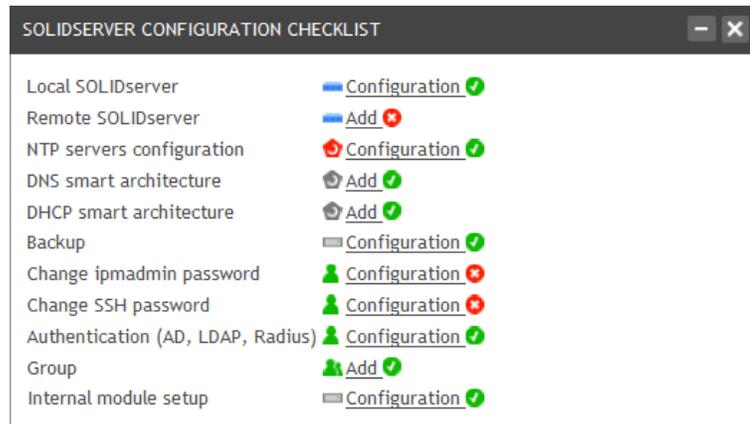


Figure 15.10. The gadget SOLIDserver configuration checklist

This gadget provides a set of shortcuts to assist you in setting SOLIDserver main configurations and making sure that your appliance is used at the best of its potential from the first connexion onward. Any line marked ✗ is not configured yet, the completed configurations are marked ✓.

The gadget provides a checklist and shortcuts toward specific configuration wizards:

Local SOLIDserver

Allows to configure locally the appliance from the gadget. Click on **Configuration** to open the wizard **Configure local SOLIDserver**. For more details, refer to the section [Configuring SOLIDserver to Remotely Manage Other Appliances](#).

Remote SOLIDserver

Allows to add remote appliances to the page *Centralized Management* from the gadget. Click on **Add** to open the wizard **Add/modify remote SOLIDserver**. For more details, refer to the section [Adding Remote Appliances](#).

NTP servers configuration

Allows to add NTP servers from the gadget. Click on **Configuration** to open the wizard **NTP servers configuration**. For more details, refer to the section [Configuring NTP Servers](#).

DNS smart architecture

Allows to create a DNS smart architecture from the gadget. Click on **Add** to open the wizard **Add a DNS server**. For more details regarding smart architectures, refer to the section [Adding a DNS Smart Architecture](#).

DHCP smart architecture

Allows to create a DHCPv4 smart architecture from the gadget. Click on **Add** to open the wizard **Add a DHCP server**. For more details regarding smart architectures, refer to the section [Adding a DHCPv4 Smart Architecture](#).

Backup

Allows to archive the appliance backup on a remote server from the gadget. Click on **Configuration** to open the wizard **Archive server parameters**. For more details regarding remote FTP configuration, refer to the section [Archiving the Backup Files on an FTP or SFTP server](#).

Change ipmadmin password

Allows the superuser, *ipmadmin*, to edit their SOLIDserver connexion password from the gadget. Click on **Configuration** to open the wizard **Modify user password**. For more details, refer to the section [Changing the Session Password](#).

Change SSH password

Allows the superuser, *ipmadmin*, to edit the SSH *admin* account password from the gadget. This account is used to authenticate users who add/edit EfficientIP DNS and DHCP servers and remote appliances. If you change the password from the gadget, you must also edit the existing servers and appliances to specify the new password. Click on **Configuration** to open the wizard **Change SSH password**. For more details, refer to the sections [Editing DNS Servers](#), [Editing a DHCP Server](#) and/or to the section [Editing Remote Appliances](#).

Authentication (AD, LDAP, RADIUS)

Allows you to add one by one the three rules that configure the remote user authentication via AD, RADIUS or LDAP. Click on **Configuration** to open the wizard **Add a rule**. For more details regarding remote authentication, refer to the chapter [Managing Authentication Rules](#).

Group

Allows you to add groups of users from the gadget. Click on **Add** to open the wizard **Add a group**. For more details, refer to the section [Adding Groups of Users](#).

Internal module setup

Allows to set the modules advanced properties interaction of the appliance. Click on **Configuration** to open the wizard **Internal module setup**. For more details, refer to the section [Defining the Internal Module Setup](#).

My Account Preferences & Configuration

This shortcut gadget is available by default on the *Main Dashboard* of all the users.



Figure 15.11. The gadget *My account preferences & configuration*

This gadget provides links to help the connected user set their preferences:

Gadgets Library

This button is a link toward the page *Gadgets Library* in the module Administration.

Set language

This button opens the wizard *Change Language*. For more details, refer to the section [Account Configuration](#).

Shortcuts

This shortcut gadget is available by default on the *Main Dashboard* of *ipmadmin*.

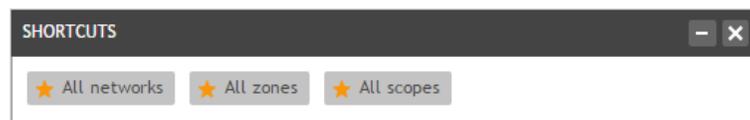


Figure 15.12. The gadget *Shortcuts*

This gadget cannot be edited and provides links to key pages of the modules IPAM, DNS and DHCP:

All networks

This button provides a shortcut toward the page *All networks* in the module IPAM. For more details, refer to the chapter [Managing Networks](#).

All scopes

This button provides a shortcut toward the page *All scopes* in the DHCP. For more details, refer to the chapter [Managing DHCP Scopes](#).

All zones

This button provides a shortcut toward the page *All zones* in the DNS. For more details, refer to the chapter [Managing DNS Zones](#).

NetChange Network Devices Vendor

This pie chart is available by default on the *NetChange dashboard* of *ipmadmin*.

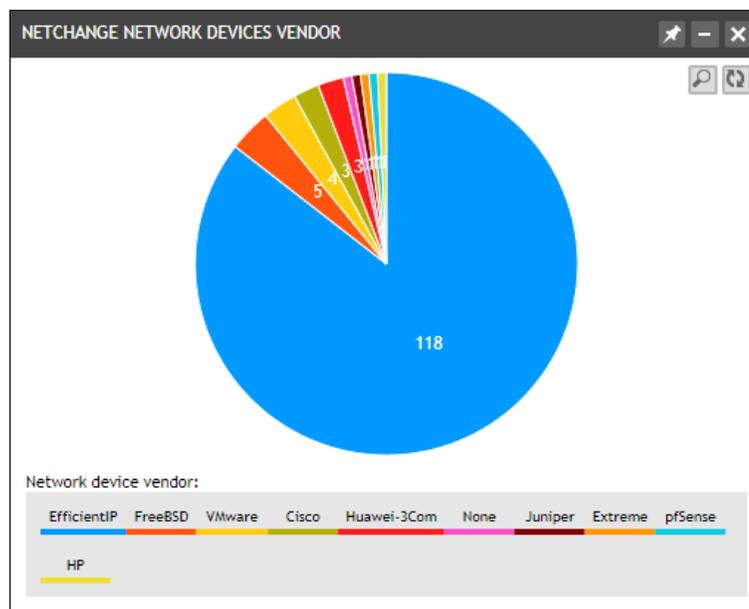


Figure 15.13. The gadget NetChange network devices vendor

This gadget represents all the vendor distribution of the ports listed on the page *All ports*, it indicates the number of ports for each by vendor. For more details, refer to the chapter [Managing Ports](#).

NetChange Active Ports Speed (bps)

This pie chart is available by default on the *NetChange dashboard* of *ipmadmin*.

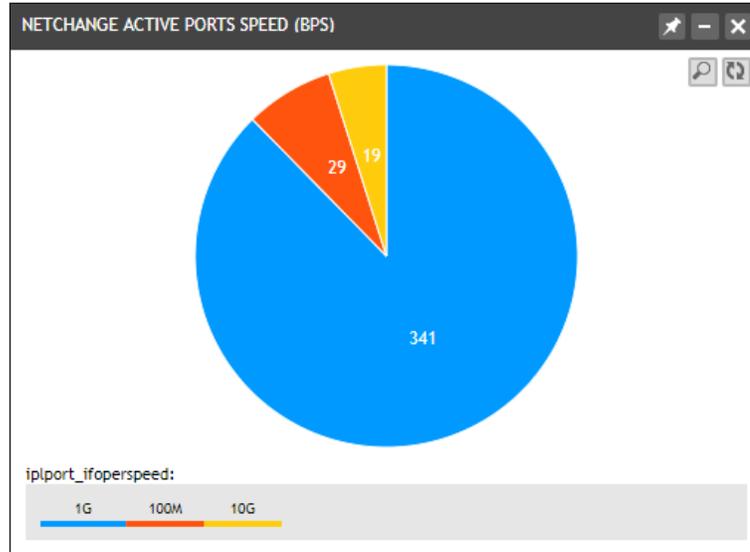


Figure 15.14. The gadget NetChange active ports speed (bps)

This gadget represents the active ports listed on the page *All ports* distributed by port speed, in bits per second, with the number of ports matching each speed. For more details, refer to the chapter [Managing Ports](#).

NetChange Port Status

This pie chart is available by default on the *NetChange dashboard* of *ipmadmin*.

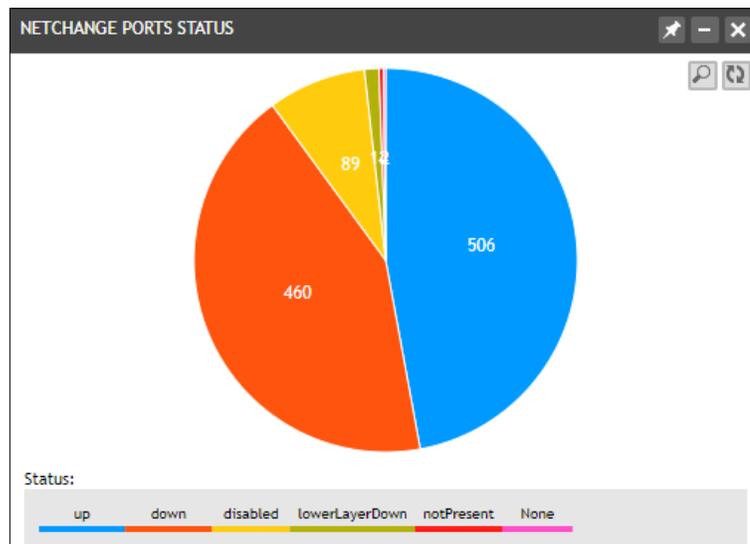


Figure 15.15. The gadget NetChange port status

This gadget represents all the ports listed on the page *All ports* distributed by status, with the number of ports marked with each status. For more details, refer to the chapter [Managing Ports](#).

Number of NetChange Ports per Device

This pie chart is available by default on the *NetChange* dashboard of *ipmadmin*.

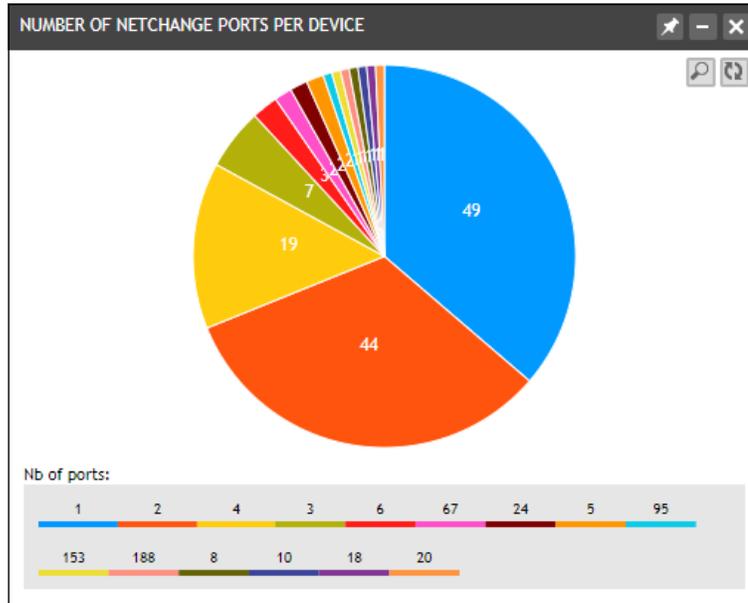


Figure 15.16. The gadget *Number of NetChange ports per device*

This gadget represents the network device repartition of all the ports listed on the page *All ports*, with the number of ports for each network device. For more details, refer to the chapter [Managing Ports](#).

Number of Ports Used per Device

This pie chart is available by default on the *Device Manager* dashboard of *ipmadmin*.

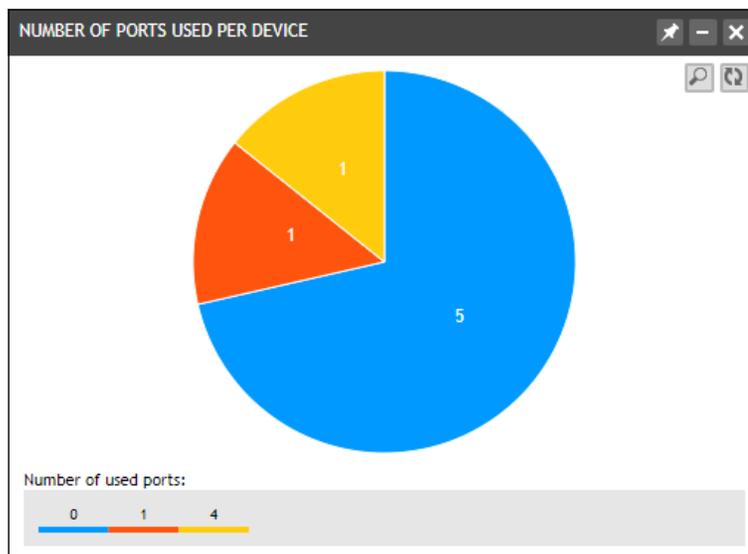


Figure 15.17. The gadget *Number of ports used per device*

This pie chart represents all the used ports listed on the page *All ports & interfaces* distributed by device. For more details, refer to the chapter [Managing Ports](#).

Number of Interfaces Used per Device

This pie chart is available by default on the *Device Manager dashboard* of *ipmadmin*.

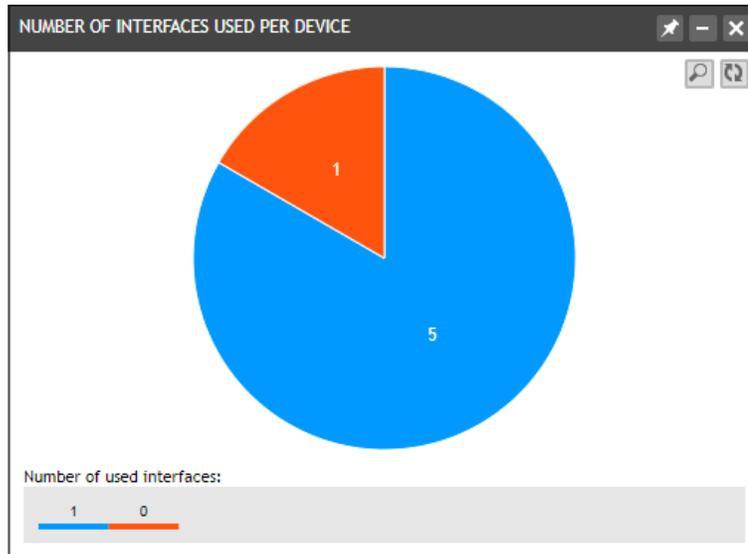


Figure 15.18. The gadget Number of Interfaces used per device

This pie chart represents all the used interfaces listed on the page *All ports & interfaces* distributed by device. For more details, refer to the chapter [Managing Ports and Interfaces](#).

Alert on Ports/Interfaces Reconciliation Drift

This Top List is available by default on the *Device Manager dashboard* of *ipmadmin*.

Starting Since	Priority	Severity
195 d 21:58	Normal	Minor
133 d 23:18	Normal	Minor
112 d 0:28	Normal	Minor
28 d 20:54	Normal	Minor
24 d 3:13	Normal	Minor

Figure 15.19. The gadget Alert on ports/interfaces reconciliation drift

This gadget provides a Top 5 list of the ports and interfaces that are marked *Drift* in the column *Reconciliation*. For more details, refer to the chapter [Managing Ports and Interfaces](#).

You can edit this Top List to display up to 25 entries in the table or edit the columns it contains. For more details, refer to the section [Editing a Top List Gadget](#).

Creating Gadgets

Each type of gadget has a specific addition and edition methods. For this reason, the creation of *Charts*, *Top List*, *Quick Search*, *Quick Wizards* and *Bookmarks* gadgets is detailed separately. Note that:

- You can create gadgets from any page of the module *IPAM*, *DHCP*, *DNS*, *Application*, *NetChange*, *Workflow*, *Device Manager* or *VLAN Manager* and some pages of the module *Administration*.
- The descriptive and configuration gadgets are default gadgets, you cannot create new ones or delete them, you can only enable or disable them and change their visibility.

Once a gadget has been created, you can assign it on any dashboard or share it with other users. For more details, refer to the sections [Assigning Gadgets to a Dashboard](#) and [Granting Other Users Access to the Gadgets](#).

Creating a Chart Gadget

You can create charts gadgets from any listing page, except in the modules *Administration* and *SPX*. Once created, you can display as many charts as you want on each dashboard.

In the GUI, some charts are already available on the page *System Statistics* or on the properties page of some resources. These charts can be assigned as gadget thanks to the button  in the drag bar. For more details, refer to the section [Assigning Gadgets to a Dashboard](#).

To create a chart gadget

1. Go to the listing page of your choice.
2. Filter the list if need be.
3. In the menu, select  - **Alerts, gadgets & smart folders** > **Add a Chart**. The wizard **Add a chart** opens.
4. Configure the chart:
 - a. In the field **Chart name**, type in the name of the gadget.
 - b. In the drop-down list **Chart type**, select *pie* or *bar*.
 - c. In the drop-down list **Value**, select the value to be displayed in the chart. On some pages, only *Entries count* is available, if so it is selected by default.
 - d. In the drop-down list **Label**, select a column. Its name is used as label.
 - e. If you selected a *bar* chart, in the drop-down list **Secondary label**, select another column. It is used to compare to the *Label*.
 - f. In the drop-down list **Order by**, you can choose to display the *Value*, *Label* or *Secondary label* in ascending order (*ASC*) or descending order (*DESC*).
5. Click on **PREVIEW** if you want to have an overview of the chart you configured.
6. Click on **NEXT**. The last page of the wizard opens.
7. In the drop-down list **Dashboard**, select the dashboard of your choice.
8. Click on **OK** to complete the operation. The report opens and closes. The chart is visible on the dashboard you chose to display it on.

Note that:

- Charts remain empty as long as there is no data to retrieve on the page they are created from.
- When you hover over a segment of a chart, a tooltip containing information about the segment is displayed.
- Chart gadgets provide specific buttons:
 - The icon  allows to zoom in on the chart in a pop-up window above the page.
 - The icon  allows to refresh the data displayed.
 - From the legend you can click on any entry to hide or display data. For more details refer to the section [Charts](#).
- Once a chart is created, you cannot edit it. If the data you configured it with no longer suits to your needs, you have to delete the gadget and create a new one. For more details, refer to the section [Deleting Gadgets](#).

Creating a Top List Gadget

You can create Top Lists gadgets from most listing page and display as many of them as you want on each dashboard.

Keep in mind that in the module Administration, you can only create a gadget Top List gadgets from the pages *Session tracking*, *User tracking* and *Alerts*.

To create a Top List gadget

1. Go to the listing page of your choice.
2. Filter the list if need be.
3. In the menu, select  **Alerts, gadgets & smart folders > Add a Top List**. The wizard **Add a Top List** opens.
4. Configure the Top List:
 - a. In the field **Name**, type in the name of the gadget.
 - b. In the drop-down list **Dashboard**, select the dashboard of your choice.
 - c. In the list **Columns**, select a column you want to display in the gadget.
 - d. Click on . The name is moved to the list **Selected columns**. Repeat these steps for all the columns you want to include to the gadget.
 - You cannot display more than 4 columns in the gadget. If the list contains more, they are ignored.
 - To remove a column from the gadget, in the list **Selected columns**, select it and click on . The column is moved back in the list **Columns**.
 - To order the columns in the Top List, select one in the field **Selected columns** and the arrows to move it up  or down .
 - e. In the field **Limit**, specify the number of items to display in the final gadget: 5, 10, 15, 20 or 25.

5. Click on **OK** to complete the operation. The report opens and closes. The gadget is visible on the dashboard you selected. It is named *Top X list: <your-gadget-name>*, where *X* is the *Limit* you selected.

You can edit the Top Lists gadgets. For more details, refer to the section [Editing a Top List Gadget](#).

Creating a Quick Search Gadget

You can create Quick Search gadgets from most listing page and display as many of them as you want on each dashboard.

The fields available to configure the Quick Search depend on the list the gadget is set from.

To create a Quick Search gadget

1. Go to the listing page of your choice.
2. In the menu, select **Alerts, gadgets & smart folders > Add a Quick Search**. The wizard **Add a Quick Search** opens.
3. Configure the Quick Search:
 - a. In the field **Name**, type in the name of the gadget. You should specify a name that indicates the page it is created from.
 - b. In the drop-down list **Dashboard**, select the dashboard of your choice.
 - c. In the list **Columns**, select the column of your choice.
 - d. Click on **+**. The column is moved to the list **Selected columns**.
 - To remove a column from the gadget, in the list **Selected columns**, select it and click on **-**. The column is moved back in the list **Columns**.
 - To order the search fields of the gadget, select a column in the field **Selected columns** and the arrows to move it up **▲** or down **▼**.
4. Click on **OK** to complete the operation. The report opens and closes. The Quick Search is visible on the dashboard you selected.

You can edit the Quick Search gadgets. For more details, refer to the section [Editing a Quick Search Gadget](#).

Creating a Quick Wizards Gadget

You can save links toward quick wizards in a Quick wizards gadget when you add or edit them.

Note that during the addition of a quick wizard you can only indicate one access location, either a dashboard for the gadget or the menu *Quick Access*. However, you can edit the quick wizard from the *My Quick Wizards* to specify more dashboards.

For more details regarding the quick wizards themselves, refer to the section [Quick Wizards](#).

To create a link in the Quick Wizards gadget toward the Quick Wizard you are creating

1. On the wizard you want to save, click on  in the wizard drag bar. The wizard **Add a Quick Wizard** opens.
2. In the field **Name**, type in the the quick wizard name.
3. In the drop-down list **Save in**, select the dashboard of your choice.
4. In the field **Description**, you can add a description.
5. Click on  to complete the operation. The report opens and closes. The Quick Wizards gadget is now displayed on the selected module dashboard, it contains a button named after your quick wizard.

To create a link in the Quick Wizards gadget toward the Quick Wizard you are editing

1. From any page, in the top bar, select  **My account > My Quick Wizards**. The page **My Quick Wizards** opens.
2. Click on the name of a quick wizard you want to edit. The wizard **Edit: Quick Wizard** opens.
3. If need be, edit the fields **Name** and **Description**.
4. In the list **Available**, select one by one the module dashboards where the quick wizard should be included to the Quick Wizards gadget.

You can also select *Quick access menu* to add a shortcut toward the quick wizard in the menu  *Quick Access*. For more details, refer to the section [Accessing Quick Wizards](#).

5. Click on . Your selection is moved to the list **Configured**. If you want to remove an item from the list, select it and click on , the item is moved back to the list **Available**.
6. Tick or untick the box **Share with other users** according to your needs. Sharing a quick wizard makes it available for all users.
7. Click on  to complete the operation. The report opens and closes. The Quick Wizards gadget is now displayed on the dashboard of the selected modules and it contains a button named after your quick wizard.

Note that once created, you can assign the Quick Wizards gadget from a dashboard. For more details, refer to the section [Assigning a Gadget from the Dashboard](#).

You can edit the content of a Quick Wizards gadget. For more details, refer to the section [Editing a Quick Wizards Gadget](#).

Creating a Gadget Bookmarks

The gadget *Bookmarks* gathers shortcuts toward the existing bookmarks of your choice. The shortcuts it contains are the same on all the dashboards it is displayed on.

You cannot create the gadget *Bookmarks* on its own, you must create it when you are bookmarking a page.

To create a gadget Bookmarks when bookmarking a page

1. On any page, click on  in the upper-right corner of the page. The wizard **Bookmark this page** opens.

2. In the field **Name**, the bookmark is named `<module>: <page>` by default. You can edit it.
3. In the field **Bookmark Folder**, you can type in a folder name to organize your bookmarks in the menu *Bookmarks*. By default, no folder is configured, the field contains `/`.
 - a. To create a folder, type in the name of your choice. You can also create sub-folders using the character `/` and the structure `/folder/sub-folder`.
 - b. To find existing folders, type in the beginning of the folder name and click on to retrieve the matching folder full name.
4. Tick the box **Add to the gadget Bookmarks** to create the gadget *Bookmarks* and add it to the page *Gadgets Library*.
5. Tick the box **Share with the other users** to make it available to all users.
6. Click on to complete the operation. The report opens and closes. The page is now bookmarked and marked ★.

The gadget is not displayed on any dashboard yet, to assign it, refer to the section [Assigning Gadgets to a Dashboard](#).

Keep in mind that:

- You cannot assign the gadget *Bookmarks* from a dashboard if you have not ticked the box *Add to the gadget Bookmarks* for at least one bookmark.
- Every time you tick the box *Add to the gadget Bookmarks*, you edit the gadget content as you add extra shortcuts to the gadget. You cannot edit the gadget from a dashboard. To remove bookmark shortcuts from the gadget, refer to the section [Editing the Bookmarks Gadget](#).

Assigning Gadgets to a Dashboard

You can assign, i.e. display, all the existing gadgets to any dashboard from:

- The dashboard itself, for more details refer to the section [Assigning a Gadget from the Dashboard](#).
- The page *System statistics*.
- The properties page of a resource that contains charts.
- The gadget related pages *Gadgets Library* and *My Gadgets*.

To assign several gadgets to several dashboards at once, refer to the section [Assigning Gadgets from the Page Gadgets Library](#).

Assigning a Gadget from the Page System Statistics

In the Administration module, the page *System statistics* allows to assign charts as gadgets.

To assign a gadget from the page System statistics

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **System statistics**. The page **System statistics** opens.
3. Click on  **Collapse all** in the upper right corner of the page to close all the panels.
4. In the panel of the chart of your choice, click on . The wizard **Gadgets Add: Chart** opens.

5. In the list **Dashboard**, select the dashboard of your choice.
6. Click on **OK** to complete the operation. The page refreshes. The gadget is now displayed on the selected dashboard.

Assigning a Gadget from a Resource Properties Page

All the charts that are displayed on the properties page of your resources can be added to any dashboard. For instance, you can assign to a dashboard any of the statistics charts available on the properties of a NetChange port.

To assign a gadget from a resource properties page

1. Go to the properties page of the object of your choice.
2. Click on **☰ Collapse all** in the upper right corner of the page to open all the panels.
3. In the panel of the chart of your choice, click on **+**. The wizard **Gadgets Add: Chart** opens.
4. In the list **Dashboard**, select the dashboard of your choice.
5. Click on **OK** to complete the operation. The page refreshes. The gadget is now displayed on the selected dashboard.

Assigning Gadgets from the Page Gadgets Library

From the page *Gadgets Library*, you can assign one or several gadgets to one or more dashboards at once.

All the existing gadgets are listed on the page, even if they are not displayed on any dashboard yet.

To assign gadgets to one or several dashboards from the page Gadgets Library

1. From any page, in the top bar, select **👤 My account > My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Tick the gadget(s) you want to assign to a dashboard.
4. In the menu, select **🔗 Edit > Assign Gadget(s)**. The wizard **Gadget configuration** opens.
5. In the list **Available**, double-click on the name of the dashboard you want the gadget to be displayed on. The name is moved to the list **Configured**. You can select several dashboards.
6. Click on **OK** to complete the operation. The report opens and closes. The gadget is now displayed on the selected dashboard(s).

Assigning a Gadget from the Page My Gadgets

The page *My Gadgets* contains only the gadgets that are already displayed on at least one dashboard. Each gadget listed can be displayed on one or more dashboards via the option *Assign a gadget*.

To assign a gadget to a dashboard from the page My Gadgets

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the menu, select  **Edit** > **Assign a gadget**. The wizard **Gadget Configuration** opens.
3. Select the **Type** that suits your needs. *Other* contains the descriptive gadgets, the gadget *Bookmarks* and the *Quick Wizards* gadget.
4. Click on **NEXT**. The list **Gadget** displays the available gadgets of the selected type.
5. Select the gadget you want. If there is no gadget of this type yet, the list is empty.
6. Click on **NEXT**. The last page of the wizard opens.
7. In the list **Available**, double-click on the name of the dashboard of your choice. The name is moved to the list **Configured**. You can select several dashboards if you want.
8. Click on **OK** to complete the operation. The report opens and closes. The gadget is now displayed on the selected dashboard(s).

Displaying or Hiding a Gadget from the Page My Gadgets

From the page *My Gadgets*, you can choose to display or hide a gadget on a particular dashboard.

To assign a gadget directly from a dashboard, refer to the section [Assigning a Gadget from the Dashboard](#).

To hide a gadget from the page My Gadgets

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. Filter the list if need be.
3. For the gadget you want to hide, in the column **Status**, click on **Visible**. The wizard **Disable Gadget** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The gadget is marked as *Hidden* and no longer visible on the dashboard.

To display a gadget from the page My Gadgets

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. Filter the list if need be.
3. For the gadget you want to display, in the column **Status**, click on **Hidden**. The wizard **Enable Gadget** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The gadget is marked as *Visible* and displayed on the dashboard it was assigned to.

If you want to hide a gadget from several dashboards at once, refer to the section [Enabling or Disabling Gadgets](#).

Editing Gadgets

You can edit some gadgets:

- The gadgets **Top list**, **Quick Search** and **Quick Wizards** can be edited from the dashboard they are displayed on. The gadgets *Top List* and *Quick Search* can also be edited from their properties page.
- The gadget **Bookmarks** can only be edited from the page *My Bookmarks*.

Any other type of gadget cannot be edited. You have to create a new gadget again and delete the one you no longer need.

Editing a Top List Gadget

You can edit a Top List gadget from one of the dashboards it is displayed on or from its properties page. Note that:

- You cannot assign a Top List gadget to a different dashboard when you edit it. To assign it to a different or more dashboards, refer to the section [Assigning Gadgets to a Dashboard](#).
- Any changes performed on a Top List gadget apply to all the dashboards it is displayed on.

To edit a Top List from a dashboard

1. Go to the dashboard of your choice.
2. In the drag bar of the Top List gadget of your choice, click on . The wizard **Edit a Top List** opens.
3. Edit the fields **Name**, **Columns**, **Selected columns** and **Limit** according to your needs.
4. Click on to complete the operation. The report opens and closes. The gadget content is updated.

To edit a Top List from its properties page

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Filter the column **Type** to only display *Top Lists*.
4. At the end of the line of the gadget of your choice, click on . The properties page opens.
5. In the panel **Main properties**, click on . The wizard **Edit a Top List** opens.
6. Edit the fields **Name**, **Columns**, **Selected columns** and **Limit** according to your needs.
7. Click on to complete the operation. The report opens and closes. The gadget content is now updated.

Editing a Quick Search Gadget

You can edit a Quick Search gadget either from one of the dashboards it is displayed on or from its properties page.

Any changes performed on a Quick Search gadget apply to all the dashboards it is displayed on.

To edit a Quick Search from a dashboard

1. Go to the dashboard of your choice.
2. In the drag bar of the Quick Search gadget of your choice, click on . The wizard **Edit a Quick Search** opens.
3. Edit the fields **Name**, **Dashboard**, **Columns** and **Selected columns** according to your needs.
4. Click on to complete the operation. The report opens and closes. The gadget content is now updated.

To edit a Quick Search from its properties page

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Filter the column **Type** to only display Quick Search gadgets.
4. At the end of the line of the gadget of your choice, click on . The properties page opens.
5. In the panel **Main properties**, click on . The wizard **Edit a Quick Search** opens.
6. Edit the fields **Name**, **Dashboard**, **Columns** and **Selected columns** according to your needs.
7. Click on to complete the operation. The report opens and closes. The gadget content is now updated.

Editing a Quick Wizards Gadget

You can edit the content of a Quick Wizards gadget from one of the dashboards it is displayed on.

You can make unique changes to a Quick Wizards gadget. Any change performed on the gadget only apply to the dashboard where you edited the gadget.

To edit a Quick Wizards gadget from a dashboard

1. Go to the dashboard of your choice.
2. In the drag bar of the Quick Wizards gadget of your choice, click on . The wizard **Edit a Quick Wizards gadget** opens.
3. Edit the gadget content:
 - a. The list **Available** contains all the existing quick wizards that are not yet displayed in the gadget.
 - b. Select the quick wizards one by one and click on  to move them to the list **Configured**. All the quick wizards listed in this field are displayed in the gadget.

- c. Repeat these actions for as many quick wizards as needed. To remove quick wizards from the gadget, select them in the list **Configured** and click on .
4. Edit the gadget content:
 - a. In the list **Available**, select the quick wizard of your choice.
 - b. Click on . The quick wizard is moved to the list **Configured**.
 - c. To remove a quick wizard from the gadget, in the list **Configured**, select it and click on . The quick wizard is moved back in the list **Available** and no longer displayed in the gadget.
5. Click on to complete the operation. The report opens and closes. The changes are only visible in the Quick Wizards gadget of the current dashboard.

Editing the Gadget Bookmarks

You can edit the content of the gadget *Bookmarks* from the page *My Bookmarks*. That is to say add or remove bookmarks from the gadget.

Any changes performed on the gadget *Bookmarks* apply to all the dashboards it is displayed on.

To add a bookmark to the gadget Bookmarks

1. From any page, in the top bar, select  **My account > My Bookmarks**. The page **My Bookmarks** opens.
2. Click on the **Name** of the bookmark you want to add to the gadget. The wizard **Edit a Bookmark** opens.
3. Tick the box **Add to the gadget Bookmarks**. The bookmark is added to the gadget *Bookmarks*.
4. Click on to complete the operation. The report opens and closes.

To remove a bookmark from the gadget Bookmarks

1. From any page, in the top bar, select  **My account > My Bookmarks**. The page **My Bookmarks** opens.
2. Click on the name of the bookmark you want to remove. The wizard **Edit a Bookmark** opens.
3. Untick the box **Add to the gadget Bookmarks**. The bookmark is removed from the gadget *Bookmarks*.
4. Click on to complete the operation. The report opens and closes.

Granting Other Users Access to the Gadgets

Only users of the group *admin* can grant access to the gadgets to other users or manage the visibility of the gadgets they created.

Standard users can create, edit or delete their gadgets, if their group is granted the proper permissions, but they cannot set the visibility of their gadgets.

Configuring Read-Write Access to a Group

Users can only manage the gadgets if the group of users they belong to is granted the relevant rights.

To grant access to the gadget related options to a group

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Groups**. The page **Groups** opens.
3. Click on the **Name** of the group of your choice. The page **Resources** opens.
4. In the breadcrumb, click on **Rights**. The page **Rights** opens.
5. In the search engine of the column **Right**, type in *gadget* to only display the gadget related rights.
6. Tick the rights that suite your needs.
7. In the menu, select  **Edit > Allow**. The wizard **Enable** opens.
8. Click on  to complete the operation. The report opens and closes. The page is visible again.

Do not forget to include users to the group. For more details regarding user permissions refer to the part [Rights Management](#).

Setting the Gadgets Visibility

On the page *Gadgets Library*, the column **All users** indicates if a gadget is visible to every user or not. Keep in mind that:

- All the users of the group *admin* can make gadgets visible to all users or only to them.
- The superuser *ipmadmin* can see all the gadgets no matter who created them or their visibility configuration.
- The gadgets displayed by default on the *Device Manager dashboard* and *NetChange dashboard* are visible to anyone with read-write access. For more details, refer to the section [Gadgets Displayed by Default](#).
- The gadgets displayed by default on the *Main Dashboard* are only visible for *ipmadmin*, unless they make them visible to all users. For more details refer to the section [Gadgets Displayed by Default](#).
- You cannot change the visibility settings of the descriptive gadgets *General information* and *System information*. They are visible by default to every user with access to gadgets.

To make a gadget visible to all users

Only users of the group *admin* can perform this operation.

1. From any page, in the top bar, select  **My account > My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.
2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.

3. Tick the gadget(s) you want to make visible to all the other users.
4. In the menu, select  **Edit** > **Visible to all users** > **Yes**. The wizard **Gadget visibility** opens.
5. Click on to complete the operation. The report opens and closes. In the column **All Users**, the gadget is marked **Yes**.

To make a gadget visible only to you

Only users of the group *admin* can perform this operation.

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Tick the gadget(s) you want to make visible only to you.
4. In the menu, select  **Edit** > **Visible to all users** > **No**. The wizard **Gadget visibility** opens.
5. Click on to complete the operation. The report opens and closes. In the column **All Users**, the gadget is marked **No**.

Enabling or Disabling Gadgets

Enabling or disabling a gadget allows to hide or display a gadget, no matter on how many dashboards it is displayed on:

- If you disable a gadget, you remove it from all the dashboards it is displayed on at once. It does not delete it. The gadget is still listed on the pages *Gadgets Library* and *My Gadgets* but its *Status* is **Disabled**.
- If you enable a gadget you previously disabled, you make it visible again on all the dashboards it is assigned to.

You can only enable or disable a gadget from the page *Gadgets Library* from the listing itself or the menu.

To enable/disable a gadget from the list

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Filter the list if need be.
4. In the column **Status**:
 - a. To enable a gadget, click on **Disabled**. The wizard **Enable Gadget** opens.
 - b. To disable a gadget, click on **Enabled**. The wizard **Disable Gadget** opens.
5. Click on to complete the operation. The report opens and closes. The gadget is marked **Enabled** or **Disabled**.

To enable/disable gadgets through the menu

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Filter the list if need be.
4. Tick the gadget(s) your choice.
5. In the menu:
 - a. To enable the gadget(s), select  **Edit** > **Status** > **Enable**. The wizard opens.
 - b. To disable the gadget(s), select  **Edit** > **Status** > **Disable**. The wizard opens.
6. Click on to complete the operation. The report opens and closes. The gadget is marked **Enabled** or **Disabled**.

Deleting Gadgets

You can only delete gadgets from the page *Gadgets Library*. Deleting a gadget removes it from the appliance altogether and no matter on how many dashboards it is displayed.

To delete a gadget

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.

Each gadget is listed one or more times, depending on how many dashboards it is assigned.

2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. Tick the gadget(s) of your choice.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The gadget is now removed from the dashboard(s) it was displayed on and from the pages *Gadgets Library* and *My Gadgets*.

Part V. IPAM

The Internet Protocol Address Management (IPAM) was designed to plan, track, organize and manage IP addresses into networks. This organization can rely on IPv4 that manages 32-bit addresses or on IPv6 that manages 128-bit addresses.

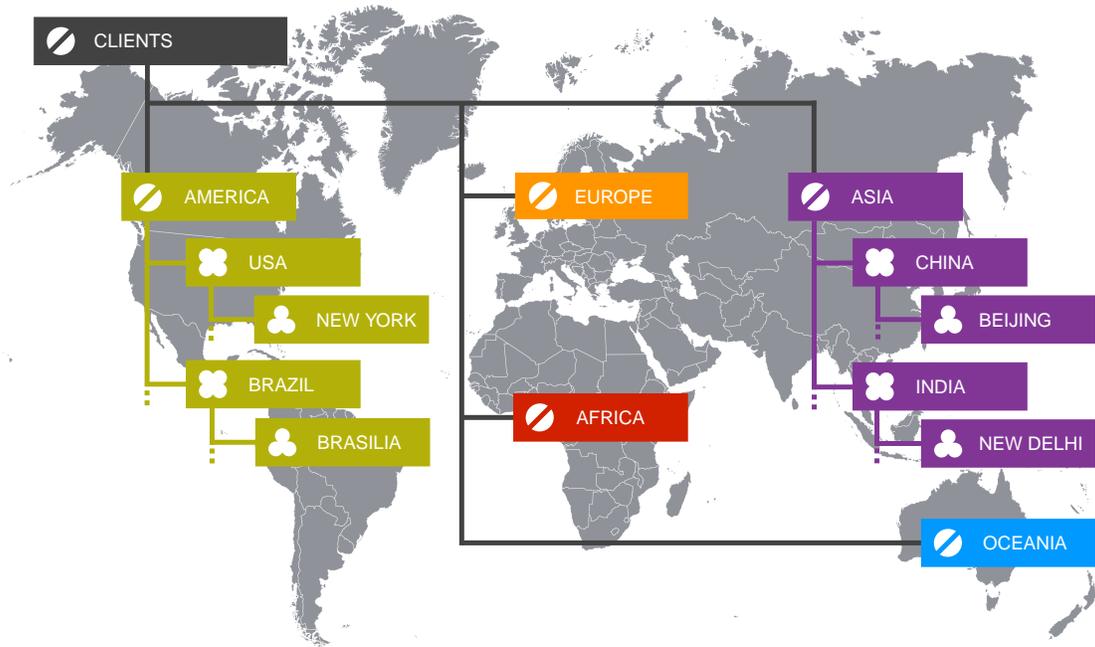


Figure 96. Representation of the IPAM hierarchy

The IP address management requires to create a space, within which you create a block-type network, that contains at least one subnet-type network that manages your IP addresses. You can also add a pool to configure some or all of your IP addresses with extra parameters. If you want to import an existing IPAM organization, refer to the part [Imports and Exports](#).

The IPAM hierarchy can include up to 5 levels of organization:

- **Space:** the highest level of the hierarchy, the essential entry point of the IP address management. It is required to create at least one. It defines the addressing space in which all the IP addresses are unique and can contain both IPv4 and IPv6 address organizations. You can create as many spaces as you need, they contain block-type networks. For more details, refer to the chapter [Managing Spaces](#).
 - **Block-type network:** the second level of the hierarchy, where you set the range of IPv4 or IPv6 addresses that you manage within your space. You must create at least one to manage IP addresses. You can create as many block-type networks as you need within a space, as long as they do not overlap each other. They contain subnet-type networks. For more details, refer to the chapter [Managing Networks](#).
 - **Subnet-type network:** the third level of the hierarchy, where you can assign IPv4 or IPv6 addresses. It is required to create at least one to manage IP addresses. You can create as many subnet-type networks as you need within a block-type network, as long as they do not overlap each other. They contain pools and/or IP addresses. For more details, refer to the chapter [Managing Networks](#).
 - **Pool:** the pool is an optional fourth level of the hierarchy. It can contain IPv4 or IPv6 addresses and allows to set them with specific parameters. For more details, refer to the chapter [Managing Pools](#).
-

-
- **IP address:** the lowest level of the hierarchy. They can belong to IPv4 or IPv6 pools and/or subnet-type networks. For more details, refer to the chapter [Managing IP Addresses](#).

Note that you can customize your organization display in the Tree view. For more details, refer to the section [Tree View](#) of the chapter [Understanding the GUI](#).

The IPAM module also provides:

- **Transition from IPv4 to IPv6.** You can configure the IPAM to automate the creation of IPv6 objects when you create them in IPv4. For more details, refer to the chapter [Setting Up a Transition From IPv4 to IPv6](#).
- **Templates for IPv4 organizations.** You can create templates to ease up the creation of IPv4 objects with specific configurations or content. For more details, refer to the chapter [Managing IPAM Templates](#).
- **Variable Length Subnet Mask (VLSM).** You can delegate and organize on different levels your spaces and/or networks in IPv4 or IPv6. For more details, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).
- **Network to VLAN association.** You can create a subnet-type network and associate it with a VLAN of the module VLAN Manager. For more details, refer to the chapter [Managing the IPAM/VLAN Interaction](#).
- **Automation of DHCP and DNS resources creation.** The advanced properties allow to automate creations in the DHCP or DNS when you create IPAM resources. For more details, refer to the chapter [Managing Advanced Properties](#).
- **SPX networks management.** Thanks to a dedicated license you can configure the connection to the RIPE or APNIC and manage your networks based on the elements you configured in the module SPX. For more details, refer to the part [SPX](#).

Note that from the module **Dashboards**, you can gather gadgets and charts on *IPAM dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 16. Managing Spaces

The space is the highest level in the IPAM module's organization, the entry point of any IPv4 or IPv6 addressing plan. It allows to manage unique ranges of IP addresses.

Browsing Spaces

Spaces are managed on the page *All spaces*. They contain the networks, pools and IP addresses.

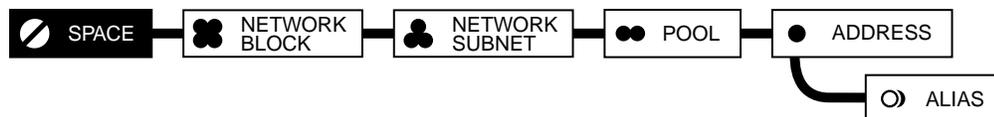


Figure 16.1. The space in the IPAM hierarchy

Spaces provide uniformity and consistency check that ensure uniqueness of IP resources: there cannot be two identical configurations of IP addresses, pools or networks within one space. To manage identical *N* address plans, you can create *N* spaces in the IPAM module.

You can create as many spaces as you want to organize your addressing plan(s) or set up multiple private networks following RFC 1918. Each space can contain as many block-type networks as you need, their size defines the number subnet-type network, pools and IP addresses that you actually manage.

By default, the space *Local* is present on the page *All spaces*. It is configured to receive all the DHCP and DNS resources configured with replication that are not attached to any space.

Browsing the Spaces Database

To display the list of spaces

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. You can filter the list using the column search engines.

To display a space properties page

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. Filter the list if need be.
3. At the end of the line of the space of your choice, click on **⌘**. The properties page opens.

Customizing the Display on the Page All Spaces

Users of the group *admin* can create customized column layouts. The button **⌘ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding a Space

Spaces allow to set up addressing plans containing either IPv4 or IPv6 addresses. You can create as many spaces as you need or use the space *Local*, created by default.

Note that you can also import spaces, for more details refer to the section [Importing Spaces](#).

To add spaces, you can use the menu or click on the button  in the upper-right corner of the page *All spaces*.

If you plan on adding different spaces with similar properties, creating a template might be useful. For more details, refer to the chapter [Managing IPAM Templates](#).

To add a space

1. In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
2. In the menu, click on  **Add**. The wizard **Add a space** opens.
3. In the list **VLSM parent space**, select *None* or one of the existing empty spaces. If you select an existing space as VLSM parent space, the new space is affiliated to the space you selected. For more details, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).
4. Click on **NEXT**. The next page opens.
5. If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Space name**, name the space.
7. In the field **Description**, you can type in a description of the space.
8. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 16.1. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

9. Click on **OK** to complete the operation. The report opens and closes. The new space is listed.

Editing a Space

At any time you can edit an existing space. You can either edit it from its properties page or via the contextual menu on the page *All spaces*.

To edit a space

1. In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
2. At the end of the line of the space of your choice, click on . The space properties pages opens.
3. In the panel **Main properties**, click on **[EDIT]**.
4. The wizard **Edit a space** opens.
5. In the list **VLSM parent space**, select a parent space if need be.
6. Click on **[NEXT]**. The next page of the wizard opens.
7. If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. Click on **[NEXT]**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

8. Edit the **Space name**, **Description** and **Advanced properties** fields according to your needs.
9. Click on **[OK]** to complete the operation. The report opens and closes. The changes are listed in the panel.

Deleting a Space

At any time you can delete a space. Keep in mind that:

- **Deleting a space also deletes all the addresses, pools and networks it contains.**
- You cannot delete a space containing other spaces. For more details regarding Variable Length Subnet Masking, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).
- None of the resources created in the DNS and DHCP from the IPAM, through the advanced properties, are deleted. This is a safety measure in case a space is deleted by mistake.

To delete a space

1. In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
2. Tick the space(s) you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on **[OK]** to complete the operation. The report opens and closes. Selected spaces are no longer listed.

Defining a Space as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a space as one of the resources of a specific group allows the users of that group to manage the space in question as long as they have the corresponding rights granted.

Granting access to a space as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 17. Managing Networks

Within the IPAM hierarchy the networks are a key level where you define ranges of IP addresses to work with. Their management follows the recommendations introduced with the RFC 950, that was aimed at providing a solution to the problems that the Internet community was facing with dual hierarchical address levels.

The networks help you set the organization that suits your needs: it can allow to set a range of addresses dedicated to your customers, and within that range a network for the customers of specific country/city; or it can help delegate terminal network management tasks to administrators.

You can manage IPv4 and IPv6 networks, within one space. To successfully set up your addressing organization, you must create within a space:

1. **A block-type network**, where you set the range of addresses you want to manage. This network is by essence non-terminal. That network contains:
2. **One or several subnet-type networks**. These networks can be terminal so you can assign the IP addresses they contain, or non-terminal and contain other subnet-type networks. For more details, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).

At network level, the IP address management follows some basic rules:

- All block-type networks belong to a space.
- Block-type networks cannot overlap each other in a space.
- All subnet-type networks belong to a block-type network.
- Subnet-type networks cannot overlap each other in a block-type network.
- A subnet-type network is defined by an IP address, size and name.

If you want to manage RIPE or APNIC networks and objects, refer to the part [SPX](#).

Browsing Networks

The page *All networks* manages both block-type and subnet-type networks, these subnet-type networks can be terminal or not.

Block-type networks, or level 0 networks, belong to spaces and are the second level of the IPAM hierarchy. They set the range of IPv4 or IPv6 addresses that you can divide into subnet-type networks. On the page they are preceded by **■**.

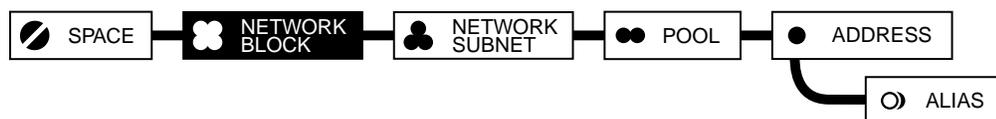


Figure 17.1. The block-type network in the IPAM hierarchy

Subnet-type networks, or level 1 (to n) networks, belong to block-type networks¹ and are the third level of the IPAM hierarchy. Terminal networks contain IPv4 or IPv6 addresses that you can assign. On the page they are preceded by ♣.



Figure 17.2. The subnet-type network in the IPAM hierarchy

You can divide the IP addresses of a subnet-type network into pools. For more details, refer to the chapter [Managing Pools](#).

The icon color provides information on the network. When ♣ is blue it indicates small sized network managing 2 or 1 IP address. It precedes /31 and /32 networks in IPv4, and /127 and /128 networks in IPv6.

Browsing the Networks Database

To display the list of networks

1. In the sidebar, go to ♣ **IPAM > Networks**. The page **All networks** opens and displays both block-type and subnet-type networks.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To display the list of networks of a specific space, in the column **Space**, click on the name of your choice. Only the networks of that space are listed.
4. To display the networks of a specific block-type network, in the column **Container**, click on the name of your choice. Only the networks of that network are listed.

In the column **Container**, *N/A* is displayed on the line of block-type networks.

To display a network properties page

1. In the sidebar, go to ♣ **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the network of your choice, click on **ⓘ**. The properties pages opens.

If you or your administrator configured IPAM to DHCP advanced properties, some subnet-type networks have a panel **DHCP options** on their properties page, to configure DHCP options for the scope associated with the network. For more details, refer to the chapters [Managing Advanced Properties](#) and [Setting DHCP Options](#).

¹Subnet-type networks can also belong to non-terminal subnet-type networks. For more details, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).

Customizing the Display on the Page All Networks

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

The page also provides columns specific to the management of SPX networks, whether RIPE or APNIC, such as *Waiting state* or *Assigned networks*. For more details, refer to the part [SPX](#).

Keep in mind that:

- In IPv4 you can display the column **Free IP**. It indicates the total number of free addresses in each subnet-type network.
- In IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the Network Statuses

The column **Status** provides information regarding the terminal networks you manage.

Table 17.1. Terminal subnet-type network statuses

Status	Description
 <i>Unmanaged</i>	The network is not managed.
 <i>OK</i>	The network is configured and managed.
 <i>Creating</i>	The delayed status while you wait for the RIPE or APNIC to confirm the network creation.
 <i>Deleting</i>	The delayed status while you wait for the RIPE or APNIC to confirm the network deletion.
 <i>NOT VALID</i>	The subnet-type network size does not fit in the block-type network although it was validated by the RIPE or APNIC. For more details, refer to the part SPX .

Adding Networks

To manage IP addresses, you must define a range of IP addresses to work with, with a block-type network, and then set the range of addresses you can assign, with a terminal subnet-type network.

To add a block-type network, refer to the section [Adding Networks Manually](#).

To add a subnet-type network, refer to the section [Adding Networks Manually](#) or [Adding Networks Using the Option By Search](#). The option *By search* allows to find the first available section of free IP addresses within a space based on a network size.

Adding Networks Manually

You can manually add IPv4 or IPv6 networks from the page All networks:

1. First, you must add a level 0 network, a block-type network, to define the range of addresses
2. Second, you can add a terminal subnet-type network. If you add non-terminal networks, you need at least one terminal network where you can assign IP addresses.

By default, the first and last IP address of a terminal network you create are reserved for the network and broadcast. The networks managing two or fewer addresses do not reserve any IP address.

Before creating a network keep in mind that:

- Several networks can be named the same.
- Several block-type networks cannot overlap each other in one space.
- Several subnet-type networks cannot overlap each other in one non-terminal network.
- By default in IPv6, you can only manually create /64, /127 or /128 terminal networks. If you want to configure them with a different prefix, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).
- You can also use templates to add several networks with similar properties. For more details, refer to the chapter [Managing IPAM Templates](#).

Note that you can also import networks, for more details refer to the section [Importing Networks](#).

To add a block-type network manually

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard opens.
4. In the drop-down list **Network type**, select *Block*². Click on **NEXT**. The next page of the wizard opens.
5. In the list **Choose a space**, select the space in which you want to add the network. Click on **NEXT**. The next page of the wizard opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Add an IPv4 network** or **Add an IPv6 network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **Network Name**, name the network.
8. In the field **Description**, you can type in a description.
9. In the field **Address**, type in the start address.
10. If you are adding an IPv4 network:
 - a. In the drop-down list **Netmask** select a netmask. The netmask value automatically edits the *Prefix*.
 - b. In the drop-down list **Prefix**, select a value if you did not choose a netmask. The prefix value automatically edits the *Netmask*.

The network size configuration is visible in the field **Comment**.

²If your group's permissions do not include the addition of both block-type and subnet-type networks, the page is automatically skipped.

11. If you are adding an IPv6 network, in the drop-down list **Prefix**, select a value between */16* and */64*. The values depend on the *Address* you specified.

If your administrator disabled the RFC 4291 compliance registry database entry, you can select a prefix between */16* and */128*. For more details, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

12. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 17.2. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

13. Click on to complete the operation. The report opens and closes. The network is listed.

Once you created a block-type network, you can create the subnet-type networks it contains.

To add a subnet-type network manually

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard opens.
4. In the drop-down list **Network type**, select *Subnet*³. Click on . The next page of the wizard opens.
5. In the list **Choose a parent space**, select a non-terminal network among the ones listed under each space. The + sign left of the spaces' name opens the list of their networks. Click on . The next page of the wizard opens.
6. You can tick the box **Allow network reparenting**. For more details, refer to the section [Reparenting subnet-type networks](#) in the chapter *Using VLSM to Manage Your IPAM Network*.
7. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on . The page **Add an IPv4 network** or **Add an IPv6 network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

8. In the field **Network Name**, name the network.
9. In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description

³If your group's permissions do not include the addition of both block-type and subnet-type networks, the page is automatically skipped.

and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).

10. In the field **Address**, type in the start address. By default, the start address of the block-type network you selected is displayed in the field.
11. If you are adding an IPv4 network:
 - a. In the drop-down list **Netmask**, select a netmask. The netmask value automatically edits the *Prefix*.
 - b. In the drop-down list **Prefix**, select a value if you did not choose a netmask. The prefix value automatically edits the *Netmask*.

The network size configuration is visible in the field **Comment**.

12. If you are adding an IPv6 network, in the drop-down list **Prefix**, select */64*, */127* or */128*.

If your administrator disabled the RFC 4291 compliance registry database entry, you can select a prefix between */16* and */128*. For more details, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

13. The box **Terminal network** is ticked by default to create a terminal network that automatically has a gateway. Depending on your administrator's display configuration:
 - a. The field **Gateway** can be displayed and you can edit it.
 - b. The field **Gateway** can be hidden but it is created anyway, based on the gateway offset calculation set in the wizard *Advanced properties customization* at network level.

You can untick the box if you want to create a non-terminal network that can contain other networks, in this case the network does not have a *Gateway*. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#).

14. If you want to add pools in this subnet network, in the field **Number of pools**, select a number.
By default, 0 is selected. If you leave this value, go to step 15.

If you select a number between 1 and 5, for each pool, two fields, **Size** and **Name**, appear.

15. Define the pools:
 - a. In the field **Size**, enter the number of IP addresses that you want in your pool.
 - b. In the field **Name**, name your pool.
 - c. Repeat the steps a and b for each pool.
16. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 17.3. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .

Field	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

17. Click on to complete the operation. The report opens and closes. The network is listed.

Adding Networks Using the Option By Search

You can use the option *By search* to create subnet-type IPv4 or IPv6 networks based on a network size. This option is designed to find matching ranges of available IP addresses within a space that you can choose from.

Before creating a network using the option *By search* keep in mind that:

- Several networks can be named the same.
- The option only allows to create subnet-type networks, they can be terminal or non-terminal.
- The wizard offers a list of the available start addresses matching your network size criteria. These results are displayed in ascending order from the non-terminal network with the most important fragmentation to the one with the least fragmentation. The hierarchy is symbolized by stars, three stars being the most.
- Several subnet-type networks cannot overlap each other in one non-terminal network.
- By default in IPv6, you can only create */64*, */127* or */128* terminal networks. If you want to configure them with a different prefix, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

To create an IPv4 network using the option *By search*

1. In the sidebar, go to  **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. In the menu, click on  **Add an IPv4 network (subnet) by search**. The wizard opens.
4. In the list **Choose a space**, select the space of your choice. Click on .
5. If you or your administrator applied classes on some block-type network(s), in the list **Parent network class** you can select a class. Selecting a class or *No class* narrows the search for available sections of IP addresses within the networks using the class selected or no class. Click on . The next page opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on . The page **Network size** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. Select a **Size**, **Prefix** or **Netmask** for your network. Selecting one value automatically changes the other two. Click on . The page **Search result** opens.

8. In the list **Network address**, select a start address. Click on **[NEXT]**. The page **Add an IPv4 subnet** opens.
9. In the field **Network Name**, name the network.
10. In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).
11. The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
12. The box **Terminal network** is ticked by default to create a terminal network that automatically has a gateway. Depending on your administrator's display configuration:
 - a. The field **Gateway** can be displayed and you can edit it.
 - b. The field **Gateway** can be hidden but it is created anyway, based on the gateway offset calculation set in the wizard *Advanced properties customization* at network level.

You can untick the box if you want to create a non-terminal network that can contain other networks, in this case the network does not have a *Gateway*. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#).

13. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 17.4. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

14. Click on **[OK]** to complete the operation. The report opens and closes. The network is listed.

To create an IPv6 network using the option **By search**

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the menu, click on **+ Add an IPv6 network (subnet) by search**. The wizard opens.
4. In the list **Choose a space**, select the space of your choice. Click on **[NEXT]**.
5. If you or your administrator created classes at network level, in the list **Parent network class** select a class or *None*. Click on **[NEXT]**. The page **Network size** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Network size** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. The box **Terminal network** is ticked by default. You can untick it if you want to create a non-terminal network that can contain other networks. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#).
8. In the drop-down list **Network prefix**, select the network prefix.

- a. If the box *Terminal network* is ticked, select *64 bits*, *127 bits* or *128 bits*.
- b. If the box *Terminal network* is not ticked, select a prefix between *8 bits* and *64 bits*.

If your administrator disabled the RFC 4291 compliance registry database entry, you can select a prefix between *8 bits* and *128 bits* whether the network is terminal or not. For more details, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

Click on **NEXT**. The page **Search result** opens.

9. In the list **Network address (v6)**, select a start address. Click on **NEXT**. The page **Add an IPv6 network** opens.
10. In the field **Network Name**, name the network.
11. In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).
12. The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
13. If the box *Terminal network* was ticked and depending on your administrator's display configuration:
 - a. The field **Gateway** can be displayed and you can edit it.
 - b. The field **Gateway** can be hidden but it is created anyway, based on the gateway offset calculation set in the wizard *Advanced properties customization* at network level.

14. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 17.5. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .

Field	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

15. Click on to complete the operation. The report opens and closes. The network is listed.

Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes

By default, the IPv6 terminal networks addition wizard offers only three prefixes: */64*, */127* and */128* to comply with RFC 4291.

If want to configure terminal networks with other, non-standard, prefixes, your administrator can break the compliance with RFC 4291 by enabling the registry database entry *module.ip.violate.rfc4291*. Once the registry entry is enabled, any prefix matching the terminal network start address is returned by the drop-down list *Network prefix*.

Note that enabling the registry database also modifies the content of drop-down list *Network prefix* for non-terminal networks.

To edit the registry key that enforces the compliance with RFC 4291

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *4291*. Only the entry *module.ip.violate.rfc4291* is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in *1* to enable it. By default, its value is *0*.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Editing Networks

You can edit networks from the page *All networks* via the contextual menu or from their properties page. Other network editions like splitting, merging or moving are detailed in the following sections.

Before editing a network properties keep in mind that you cannot edit networks created using a template. For more details, refer to the chapter [Managing IPAM Templates](#).

To edit a network

1. In the sidebar, go to  **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the network of your choice, click on . The properties pages opens.
4. In the panel **Main properties**, click on . The wizard opens.

5. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Edit an IPv4 network** or **Edit an IPv6 network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Edit the **Network name**, **Description** according to your needs.
7. If you are editing a subnet-type network, the box **Terminal network** is displayed.
 - a. You can tick the box, in this case the network is terminal and has a gateway. Depending on your administrator's display configuration, the field **Gateway** can be displayed and you can edit it.
 - b. You can untick the box, in this case the network is non-terminal and can contain other networks. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#).

Keep in mind that you cannot make a terminal network non-terminal or vice versa if they contain other networks or pools.

8. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 17.6. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

9. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Splitting Networks

You can split IPv4 or IPv6 networks to manage the IP addresses they contain in separate networks. Before splitting network, keep in mind that:

- You can split a network into 2, 4 or 8 smaller networks of same size.
- The new networks are all named after the network you split.
- When splitting a terminal network, the new smaller networks all use the first available IP address as their gateway.
- You can split a network containing other networks only if the operation does not impact the networks it contains.
- If the networks are configured with advanced properties, splitting them might compromise your original configuration.

- You cannot split *unmanaged* networks.

To split a network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the network(s) you want to split.
4. In the menu, select **Edit > Split**. The wizard **Splitting networks** opens.
5. In the drop-down list **Number of networks to create**, select 2, 4 or 8. By default, 2 is selected.
6. Click on **OK** to complete the operation. The report opens and closes. The networks are listed.

Merging Networks

You can merge IPv4 or IPv6 terminal networks, to manage more addresses. Before merging subnet-type networks, keep in mind that:

- You can only merge subnet-type networks.
- You can only merge networks belonging to the same non-terminal network, either a non-terminal subnet-type network or a block-type network.
- You can only merge contiguous networks.
- You can only merge networks of equal size.
- The new network is named after the very first network in the list that manages all the IP addresses of the selected networks.
- The number of networks before the merge must be a power of two (2, 4, 8, 16, 32, 64, ...).
- The new larger network uses one of the existing gateway addresses as gateway for the new network.
- The result of the merge must produce a network with a netmask address boundary.
- You cannot merge *unmanaged* networks, IPv6 networks or block-type networks.

To merge subnet-type networks

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the networks you want to merge.
4. In the menu, select **Edit > Merge**. The wizard **Merging networks** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The new network is listed, it is named after the first network in the list and has the start address of the first selected network and the end address of the last selected network.

Keep in mind that if the merge you are trying to execute is impossible, an error message appears on the report page and only a partial report of some networks is executed.

Moving Networks

You can move IPv4 or IPv6 networks from one space to the other. Before moving a network keep in mind that:

- The migration of a block-type network also moves the networks, pools and IP addresses it contains.
- The migration of a subnet-type network is only possible if the target space contains a block-type network that can receive it.
- You can overwrite an existing network in the target space if both networks have the same size.
- You cannot move a network if it overlaps partially an existing network in the target space.
- You cannot move an *unmanaged* network.

To move a network from one space to the other

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v4]** or **[v6]** depending on your needs. The page refreshes and the button turns black.
3. Tick the network(s) you want to move.
4. In the menu, select **Tools > Expert > Migrate to another space**. The wizard **Migrating IPv4 networks** or **Migrating IPv6 networks** opens.
5. In the drop-down list **Target space**, select the space where you want the network to be moved.
6. In the drop-down list **Overwrite**, select *Yes* or *No* according to your needs.
7. Click on **[OK]** to complete the operation. The report opens and closes. The selected network's space has changed.

Discovering the Assigned IP Addresses in a Network

Within terminal IPv4 networks, you can find out which IP addresses are being used. The option *Discover networks* allows to ping the IP addresses of the terminal network and assign them on the page *All addresses*. That way, the IP addresses already in use on your physical network - computers, servers, etc. - cannot be used. **If advanced properties are configured, this addition may also update the database of other modules.** For more details, refer to the section [IP Address Advanced Properties](#).

The option *Discover networks* names the assigned IP addresses if the DNS resolver of the appliance is properly configured and if these IP addresses are declared in a valid PTR record. For more details, refer to the sections [Setting the DNS Resolver](#) and [Adding a PTR Record](#).

Considering that pinging all the IP addresses of a network can take some time, you can choose to perform this scan at different speeds: it can be fast, normal or slow. The slower the discovery, the more likely you are to properly scan the network. The discovery mechanism sends 32 ICMP echoes at once on the network. For more details, refer to the table below.

Table 17.7. The available speeds for the option Discover networks

Speed	Timeout	Retry
Slow	3 seconds	2 attempts

Speed	Timeout	Retry
Normal	2 seconds	1 attempt
Fast	1 second	no retry

To discover the assigned IPv4 addresses within a terminal network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Tick the terminal network(s) of your choice.
4. In the menu, select **Edit > Discover networks > Fast, Normal or Slow**. The wizard **Discover networks** opens.
5. Click on **OK** to perform the terminal networks discovery. The report opens during the discovery. Click on **CLOSE** to go back to the list.

When the operation is over, the assigned IP addresses are the ones that responded to the ping. They are marked **In use** in the column *Status* and they have a *Name* if they are properly declared in a PTR record.

Using Network Map to Display Assigned IP Addresses

From the page *All networks* of a specific non-terminal network you can access the page **Network map** to have an overview of which segments are already managed and in which segments you can create new subnet-type networks, terminal or not.

Network map allows to see the occupancy rate of IPv4 block-type or subnet-type non-terminal networks and to make sure you did not forget any IP addresses in your addressing strategy.

To display the page Network map

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Click on the name of the non-terminal network of your choice. The page refreshes and displays the networks it contains.
4. In the breadcrumb on the right of the network name, click on **»** to display additional pages.
5. Click on **Network map**. The page refreshes.
6. To access the properties page of a subnet-type network, click on any blue area.

On the page Network map, non-terminal networks are divided into lines of /24 terminal networks, where free ranges are gray and used ranges are blue.



Figure 17.3. Example of a Network map

- 1 This column indicates the first IP address of the networks. Every simple line represents a segment of 256 consecutive IP addresses, free or belonging to a network.
- 2 The blue areas indicate used IP addresses within the non-terminal network. On this image, three existing networks are highlighted:
 - 3.2.12.0-3.2.15.255: This network is dark and light blue because it uses more than 256 addresses. The light area indicates that no IP address is free in the portion of the non-terminal network.
 - 3.2.50.0-3.2.50.255: This network is dark blue because it uses exactly 256 addresses.
 - 3.2.255.0-3.2.255.63: This network uses only a portion of a /24 network, so it is directly followed by a gray area that extends to the next existing network to highlight all the free IP addresses between the two networks.
- 3 The gray areas indicates the free IP addresses within the non-terminal network. On this image, the segment 3.3.4.0-3.3.255.255 is free.
- 4 Put your mouse over any blue area to display the network details: name, start and end IP addresses and size. If you click on the blue area the properties page of the network opens.
- 5 This column indicates the last IP address of the network.

Managing or Unmanaging Networks

You can choose to manage or unmanage IPv4 or IPv6 networks, unmanaging a network prevents overlapping. By default, all networks are managed, in the column *Status* they are **OK**.

The option can be useful, for instance, if you are allocated a particular range of addresses by the RIPE or APNIC through SPX, especially if you are still waiting on this range to be officially allocated to you. As any network set as *unmanaged* is virtually non-existent in the database, it gives you time to create new networks managing the same start IP address and prefix as an existing unmanaged network and assign in advance the addresses it contains if need be.

Before using the option, keep in mind that:

- You can only use the option on terminal networks, i.e. subnet-type networks.
- Unmanaging a subnet-type network puts the used IP addresses it manages in a container *Orphan addresses*.
- It is impossible to manage again an unmanaged network if it uses IP addresses that are already used by a *managed* network.
- It is impossible to split or merge unmanaged networks.

To manage/unmanage a subnet-type network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the subnet-type network(s) of your choice.
4. In the menu, select **Tools > Expert > Manage** or **Unmanage**. The wizard opens.
5. Click on **OK** to complete the operation. The report opens and closes. In the column **Status**, the selected networks are marked **OK** or **Unmanaged**.

Automating the DNS Records Update

From the page *All networks*, you can configure and apply a link between a subnet-type network and a zone via the advanced properties. This allows you to automate record creations when adding or editing IP addresses within the subnet-type networks of your choice.

- The advanced properties allow to set a default domain (zone) to associate with subnet-type networks or to set up a list of zones to choose from when adding or editing subnet-type networks. For more details, refer to the chapter [Managing Advanced Properties](#).
- Once these links are applied, the IP address creation or edition automatically creates records in the selected zone if you apply the property "Update DNS" at IP address level. For more details, refer to the chapter [Managing Advanced Properties](#).

Automating the DHCP Statics Reservation

From the page *All networks*, you can use the advanced properties to link IPAM subnet-type networks to DHCP clusters. This allows to automate static creation when adding or editing IP addresses within the subnet-type networks of your choice.

When you add IP addresses, you can automatically reserve a static in the DHCP thanks to the advanced properties.

- You can link a subnet-type network to the cluster of your choice. This allows to automate static reservations on the selected servers. For more details, refer to the chapter [Managing Advanced Properties](#).
- Once this link is set, the edition of the IP addresses within the configured subnet-type networks reserves statics using the IP address details if you apply the property "Create DHCP static" at IP address level. For more details, refer to the chapter [Managing Advanced Properties](#).

Associating Networks With a VLAN

SOLIDserver provides the possibility of creating Virtual Local Area Networks in the module VLAN Manager and associate them networks. This configuration allows two subnet-type networks to communicate with each other no matter what space or network they belong to.

You can set up this configuration between existing VLANs and subnet-type networks directly from the IPAM using the advanced properties. For more details, refer to the section [Managing the IPAM/VLAN Interaction](#) in the part VLAN Manager.

Deleting Networks

At any time you can delete networks. But keep in mind that:

- As a safety measure:
 - Deleting a block-type network does not erase it from the database. If it contains objects - networks, pools, used IP addresses - it is renamed *Orphan Network* and the objects it contained are managed by the next block-type network created that matches the deleted network configuration, i.e. the same start IP address, size and advanced properties.
 - Deleting a subnet-type network puts its used IP addresses in an *Orphan Network*. The free IP addresses are deleted.

If you really want to delete a network, you must delete its content first, starting with its used IP addresses and up the IPAM hierarchy until the network you want to delete is empty. For more details, refer to the chapters [Managing IP Addresses](#) and [Managing Pools](#).

- Deleting a network configured with the advanced property *DNS server for reverse zones* or *DNS view for reverse zones* also deletes the corresponding reverse zone from the DNS, if it only contains the default SOA and NS records.

To delete networks from a VLSM organization, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).

To delete a network

1. In the sidebar, go to  **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the network(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. Selected networks are no longer listed, they might be replaced by *Orphan networks* or *Orphan Addresses*.

Defining a Network as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a network as one of the resources of a specific group allows the users of that group to manage the network in question as long as they have the corresponding rights granted.

Granting access to a network as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 18. Managing Pools

Within the IPAM hierarchy, pools are the fourth level of the IPAM module hierarchy, they are the last container level. They can be created within terminal networks to manage IP addresses, their use is optional.

Pools allow reserving IP addresses for restricted usage such as: address provisioning, planning or migrations. Pools can also be used to delegate one or several ranges of IP addresses to groups of administrators or to restrict access to users.

Browsing Pools

The pools belong to terminal networks and contain IPv4 or IPv6 addresses. They are managed in the page *All pools*.



Figure 18.1. The pool in the IPAM hierarchy

Pools are identified by name and start/end IP address.

Browsing the Pools Database

To display the list of pools

1. In the sidebar, go to ♣ **IPAM** > **Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To display the list of pools of a specific network, in the column **Network**, click on the name of the network of your choice. The main properties of the network are displayed.
4. In the breadcrumb, click on **All pools**. The list of pools that this specific network contains is displayed.

To display a pool properties page

1. In the sidebar, go to ♣ **IPAM** > **Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the pool of your choice, click on **ⓘ**. The properties pages opens.

Customizing the Display on the Page All Pools

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that in IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Adding a Pool

Within any terminal network, a subnet-type network, you can create pools to organize further your IP addresses and, for instance, configure them with a common set of options. Pools can be created from the pages *All pools*, *All addresses* and from a terminal network's properties page, in the panel *IP address pool* of .

Note that you can also import pools, for more details refer to the section [Importing Pools](#).

To add a pool

1. In the sidebar, go to  **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Click on the name of the subnet-type network of your choice. The page **All addresses** of this network opens.
4. In the breadcrumb, click on **All pools**. The page **All pools** of the network opens.
5. In the menu, click on  **Add**. The wizard opens.
6. If you or your administrator created classes at pool level, in the list **IP pool class**, select a class or *None*. Click on . The page **Add an IPv4 pool** or **Add an IPv6 pool** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **Pool name**, name the pool.
8. The box **Pool read-only** allows to reserve the pool, mark it as read-only. By default, the box is not ticked. For more details, refer to the section [Reserving a Pool](#).
9. In the field **Start address**, type in the first address of the pool.
10. In the field **End address**, type in the last address of the pool. Specifying an end address automatically edits the value of the field *Size*.
11. In the field **Size**, you can type in the number of IP addresses you want to manage with the pool. If you specify a size, the end address is automatically modified. The value of the field *Size* is automatically calculated using the fields *Start address* and *End address*.
12. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 18.1. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The pool is listed. On the page *All addresses*, the column *Pool* indicates the pool name next to all the addresses it manages.

Reserving a Pool

You can reserve a pool when creating or editing a pool thanks to the box *Pool read-only*. This reservation may be useful to dedicate the use of IP addresses to the DHCP, identify a bunch of printers, etc.

To reserve an existing pool

- Go to the properties page of the pool of your choice.
- In the panel **Main properties**, click on . The wizard opens.
- If you or your administrator created classes at pool level, in the list **IP pool class**, select a class or *None*. Click on . The page **Edit an IPv4 pool** or **Edit an IPv6 pool** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Tick the box **Pool read-only** to reserve the pool.
- Click on to complete the operation. The report opens and closes. The pool is now marked **Yes** in the section *Read-only* of the panel **Main properties**.

When adding a pool, you can tick the box as well to reserve it.

Resizing a Pool

You can resize IPv4 pools to manage more or less addresses than they did when you created them. Resizing a pool shifts the start and/or end IP address of the pool, you can specify a number of addresses to include/exclude.

So, if your pool managed the addresses *192.168.100.10-192.168.100.125* you can decide to resize it to manage the addresses *192.168.100.100-192.168.100.105* indicating a start address shift of "90" and an end address shift of "-20".

You cannot resize a pool if the addresses you include or exclude are already used or belong to another pool.

To resize an IPv4 pool

1. In the sidebar, go to  **IPAM** > **Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the pool(s) you want to resize.
5. In the menu, select  **Edit** > **Resize Pools**. The wizard **Resize IPAM pools** opens.
6. In the field **Start address shift**, type in the positive or negative start address shift for the pool. If you type in 0 (zero), the address stays the same.
7. In the field **End address shift**, type in the positive or negative end address shift for the pool. If you type in 0 (zero), the address stays the same.
8. Click on  to complete the operation. The report opens and closes. The new start/end address for the pool(s) is listed.

Deleting a Pool

At any time, you can delete pools. Note that **if you delete a pool you do not delete the addresses it contains** or create an *orphan* container. You only delete the pool itself and the parameters it was set with.

If addresses inherited class parameters from the deleted pool, their value and the source of their value remain the same: the *Inheritance property* of each class parameter is forced to *Inherit* or *Set* to match the configuration of the deleted pool.

To delete a pool

1. In the sidebar, go to  **IPAM** > **Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Tick the pool(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on  to complete the operation. The report opens and closes. The pool is no longer listed.

On the page *All addresses*, the free IP addresses managed by a deleted pool are no longer listed. The addresses *In use* are still listed but next to them the column *Pool* is empty.

Defining a Pool as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a pool as one of the resources of a specific group allows the users of that group to manage the selected pool as long as they have the corresponding rights granted.

Granting access to a pool as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 19. Managing IP Addresses

The IP address is the last level of the IPAM hierarchy, where you assign your IP addresses to specific users, devices, etc.

The page *All addresses* provides management options for IPv4 and IPv6 and allows to display on one page the entire IP address database.

Browsing IP Addresses

The IP addresses can belong to terminal subnet-type networks and pools. They can be configured with one or several aliases.

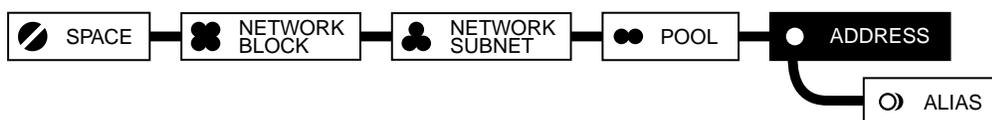


Figure 19.1. The IP address in the IPAM hierarchy

Browsing the IP Addresses Database

To display the list of addresses

1. In the sidebar, go to ♣ **IPAM** > **Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To display the list of addresses of a specific network, in the column **Network**, click on the name of the terminal network of your choice. The page refreshes, the list only displays the IP addresses of that network; the breadcrumb indicates the parent network(s) of the terminal network.

To display IPv6 addresses in full

1. In the sidebar, go to ♣ **IPAM** > **Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. On the right-end side of the menu, click on **0::0 Uncompress IPv6 addresses**. The page refreshes and all the addresses are displayed entirely.

To display an IP address properties page

1. In the sidebar, go to ♣ **IPAM** > **Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the address of your choice, click on **⌘**. The properties pages opens.

Customizing the Display on the Page All Addresses

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that:

- In IPv4 you can display the column **Aliases**. It provides a complete overview of the aliases configured on IPv4 addresses.
- In IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the IP Address Type and Status

The columns **Type** and **Status** provide information regarding the IP addresses you manage.

These columns provide information regarding the IPAM to DHCP interaction that can be set via the advanced properties or directly when configuring statics. For more details, refer to the chapter [Managing Advanced Properties](#) or the section [Adding DHCPv4 Statics](#).

Table 19.1. IP address types

Type	Description
 <i>Network</i>	The Network address is the first IP address of an IPv4 or IPv6 terminal network. It should not be assigned, except in IPv4 networks with a /31 or /32 prefix or in IPv6 networks with a /127 or /128 prefix. To make it assignable refer to the section Assigning the Broadcast and Network Addresses .
 <i>Broadcast</i>	The Broadcast address is the last IP address of an IPv4 terminal network. It should not be assigned, except in IPv4 networks with a /31 or /32 prefix. To make it assignable refer to the section Assigning the Broadcast and Network Addresses .
 <i>Orphan</i>	The IP address belonged to a network that was deleted. It was not deleted because it is assigned.
 <i>Regular</i>	The IP address is only managed in the IPAM.
 <i>DHCP static</i>	Only for IPv4, the IP address was added in the IPAM and configured with the advanced property <i>Create a DHCP static</i> . For more details, refer to the chapter Managing Advanced Properties .
 <i>DHCP lease</i>	Only for IPv4, the IP address belongs to a pool configured with the advanced property <i>Add a DHCP range</i> . For more details, refer to the chapter Managing Advanced Properties .

Table 19.2. IP address statuses

Status	Description
 <i>Non assignable</i>	The IP address cannot be assigned, it is the  <i>Network</i> or  <i>Broadcast</i> address of a terminal network. For more details, refer to the table IP address types .
 <i>Reserved</i>	Only for IPv4, the IP address is reserved for DHCP use, if it is a  <i>DHCP lease</i> it can be allocated to a DHCP client the next time they connect to the network. Only for IPv4, the IP address is reserved for DHCP use, if it is a  <i>DHCP static</i> it is assigned to a DHCP client.
 <i>In use</i>	For a  <i>Regular</i> IP address, the address is already assigned. For a  <i>DHCP lease</i> , the IP address was allocated to a DHCP client that connected to the network.

Status	Description
	For a DHCP static , the IP address was assigned to a DHCP client and allocated as a lease to the client when they connected to the network. This status is only displayed if your DHCP configuration relies on EfficientIP DHCP servers in version 6.0.0 or more.
<i>Free</i>	The IP address can be assigned.
<i>Read-only / Free</i>	The IP address is currently free but cannot be assigned because it belongs to a pool in read-only. That pool is not configured with DHCP advanced properties.
<i>Read-only / In use</i>	The IP address is currently used but cannot be edited because it belongs to a pool in read-only. That pool is not configured with DHCP advanced properties.
<i>Invalid</i>	Only for IPv4, the IP address does not match any IP address in the DHCP database but it belongs to a network or pool configured with DHCP advanced properties. The IPAM to DHCP advanced properties are probably misconfigured.

Keep in mind that the **Gateway** address of a terminal network is different from the other IP addresses:

- It can be automatically created when you create a terminal network and named *Gateway*. For more details, refer to the chapter [Configuring IPAM Advanced Properties](#).
- Its address can be automatically calculated to use the penultimate address of the network, unless you edit the value of the field *Gateway offset* in the wizard *Advanced properties customization*. For more details, refer to the chapter [Configuring IPAM Advanced Properties](#).

Adding an IP Address

At address level, there are two ways of adding, or assigning, IP addresses:

- **Manually:** if you already know the IP address you want to assign and are sure that this IP address is free. You can add it from the menu *Add* or from the list *All addresses* itself.
- **By search:** if you do not know if there is a free IP address within your terminal networks, you can use this option to find available IP addresses.

Note that you can also import IP addresses, for more details refer to the section [Importing IP Addresses](#).

When you create large terminal networks, the Broadcast and Network addresses are automatically assigned. Both addresses are by default *Non assignable* in large networks, but you can make them assignable. For more details refer to the section [Assigning the Broadcast and Network Addresses](#).

Adding an IP Address Manually

From the page *All addresses* you can add IP addresses, that is to say assign them a name and associate them with a specific MAC address. You can do it from the menu or from the list of IP addresses.

To add an IP address from the menu

1. In the sidebar, go to **IPAM** > **Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.
3. In the menu, select **+ Add** > **Address** or **Address (v6)**. The wizard opens.

4. In the list **Space**, select a space and click on **NEXT**. The next page opens.
5. If you or your administrator created classes at higher levels, the **<object> class** page appears. Select the class of your choice, *All* or *None/No class* and click on **NEXT**. The next page opens. For more details, refer to the section [Applying Classes](#).
6. In the list **Network name**, select the terminal network of your choice and click on **NEXT**. The next page opens.
7. If the network contains pools, in the list **Pool name**, select a pool or *No pool*. Click on **NEXT**.
8. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. On the page **Add an IPv4 address** or **Add an IPv6 address**, configure your IP address:

Table 19.3. IP address configuration fields

Field	Description
IP address name	This field is in read-only and displays the name specified in the field <i>Shortname</i> followed by the domain if you select one.
IP address	The IP address you assign, has to be part of the selected terminal network. By default, the first free and Regular IP address in this network or the first IP address of the selected pool is displayed in the field.
MAC address	You can type in a MAC address for the IP address. Remember that in IPv6, the MAC address corresponds to the last twelve hexadecimal characters of the client DUID. This field is optional.
Shortname	Name the IP address.
Domain	You can select one of your DNS zones or <i>None</i> . If you select a domain, the IP address updates the selected zone; its name is edited and follows the syntax: <i><shortname.domain></i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

10. Click on **NEXT**. The page **Aliases configuration** opens.
11. In the field **Add an alias**, name your alias(es). Click on **ADD** to add it to the **Aliases list**. For more details regarding aliases, refer to the section [Configuring and Managing IP Address Aliases](#).
12. Click on **OK** to complete the operation. The report opens and closes. The address is listed, its **Type** and **Status** depend on your configuration. For more details, refer to the section [Understanding the IP Address Type and Status](#).

From the page *All addresses*, whether it displays all the addresses or only the addresses of a specific network, you can click on any *free* IP address to name and configure it.

To add an IP address from the list

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v4]** or **[v6]** depending on your needs. The page refreshes and the button turns black.
3. Click on the available IP address of your choice. A pop-up window **This address is free, do you want to assign it?** opens.
4. Click on **[OK]**. The wizard opens.
5. If you or your administrator created classes at higher levels, the **<object> class** page appears. Select the class of your choice, *All* or *None/No class* and click on **[NEXT]**. The next page opens. For more details, refer to the section [Applying Classes](#).
6. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. On the page **Add an IPv4 address** or **Add an IPv6 address**, configure your IP address:

Table 19.4. IP address configuration fields

Field	Description
IP address name	This field is in read-only and displays the name specified in the field <i>Shortname</i> followed by the domain if you select one.
IP address	This field is in read-only and displays the IP address you clicked on in the list.
MAC address	You can type in a MAC address for the IP address. Remember that in IPv6, the MAC address corresponds to the last twelve hexadecimal characters of the client DUID. This field is optional.
Shortname	Name the IP address.
Domain	You can select one of your DNS zones or <i>None</i> . If you select a domain, the IP address updates the selected zone; its name is edited and follows the syntax: <i><shortname.domain></i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

8. Click on **[NEXT]**. The page **Aliases configuration** opens.
9. In the field **Add an alias**, name your alias(es). Click on **[ADD]** to add it to the **Aliases list**. For more details regarding aliases, refer to the section [Configuring and Managing IP Address Aliases](#).
10. Click on **[OK]** to complete the operation. The report opens and closes. The address is listed, its **Type** and **Status** depend on your configuration. For more details, refer to the section [Understanding the IP Address Type and Status](#).

Adding an IP Address Using the Option By search

From the page *All addresses* you can add addresses using the option *By search*. This option allows to automatically find available IP addresses within a terminal network and configure the one that suits your needs.

To add an IP address using the option By search

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on **+ Add an IPv4 address by search** or **Add an IPv6 address by search**. The wizard opens.
4. In the list **Space**, select a space and click on **NEXT**. The next page opens.
5. If you or your administrator created classes at higher levels, the **<object> class** page appears. Select the class of your choice, *All* or *None/No class* and click on **NEXT**. The next page opens. For more details, refer to the section [Applying Classes](#).
6. In the list **Network name**, select the terminal network of your choice and click on **NEXT**. The next page opens.
7. If the network contains pools, in the list **Pool name**, select a pool or *No pool*. Click on **NEXT**. The page **IP address class** opens.
8. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. In the list **IP address**, select the IP address of your choice. The first ten available addresses of the network are listed. Click on **NEXT**.
10. On the page **Add an IPv4 address** or **Add an IPv6 address**, fill in the following fields:

Table 19.5. IP address configuration fields

Field	Description
IP address name	This field is in read-only and displays the name specified in the field <i>Shortname</i> followed by the domain if you select one.
IP address	This field is in read-only and displays the IP address you selected in the field <i>IP address</i> on the previous page.
MAC address	You can type in a MAC address for the IP address. Remember that in IPv6, the MAC address corresponds to the last twelve hexadecimal characters of the client DUID. This field is optional.
Shortname	Name the IP address.
Domain	You can select one of your DNS zones or <i>None</i> . If you select a domain, the IP address updates the selected zone; its name is edited and follows the syntax: <i><shortname.domain></i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .

Field	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on . The page **Aliases configuration** opens.
- In the field **Add an alias**, name your alias(es). Click on to add it to the **Aliases list**. For more details regarding aliases refer to the section [Configuring and Managing IP Address Aliases](#).
- Click on to complete the operation. The report opens and closes. The address is listed, its **Type** and **Status** depend on your configuration. For more details, refer to the section [Understanding the IP Address Type and Status](#).

Assigning the Broadcast and Network Addresses

By default, the IPv4 terminal networks with a prefix smaller than /30 (more than 4 addresses) are created with a **Network** address and **Broadcast** address that are both **Non assignable**. In Pv6, terminal networks with a prefix smaller than /127 (more than 3 addresses) are created with a Non assignable Network address.

Depending on the way you organized your addressing plan, you might need to assign all the addresses of your networks, including the Network and Broadcast addresses. A registry database entry allows to make them both assignable.

To make the Broadcast and Network addresses of a network assignable

Only users of the group *admin* can perform this operation.

- In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- In the search engine of the column **Name**, type in *www.display.lock_broadcast_network_addresses* and hit **Enter**. The key is the only one listed.
- In the column **Value**, click on the key value. The wizard **Registry database Edit a value** opens.
- In the field **Value**, type in *0* to remove the restriction on the broadcast and network addresses. By default, the value of the key is *1*.
- Click on to complete the operation. The report opens and closes. The **Registry database** is visible again.

Once the Broadcast and Network addresses are assignable, you can configure them like any other IP address. Note that on the page *All addresses*, their **Status** remains *Non assignable* but the IP address itself is underlined to indicate that you can assign them.

Editing an IP Address

You can edit used IP addresses to change their class, name, MAC address or even their advanced properties or class parameter configuration. This edition can be done from the list *All addresses* or from the IP address properties page.

Note that you cannot edit IP addresses belonging to a pool read-only.

To edit an IP address

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Right-click over the **Name** of the IP address you want to edit. In the contextual menu, click on **Edit**. The wizard opens.
4. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the IP address configuration according to your needs:
 - a. You cannot edit the fields **IP address** and **IP address name**. *IP address name* displays the changes performed in the fields *Shortname* and/or *Domain*.
 - b. Edit the **MAC address**, **Shortname**, **Domain** and/or **Advanced properties** configuration. For more details, refer to the chapter [Managing Advanced Properties](#).
6. Click on **NEXT**. The page **Aliases configuration** opens.
7. In the field **Add an alias**, name your alias(es). Click on **ADD** to add it to the **Aliases list**. For more details regarding aliases, refer to the section [Configuring and Managing IP Address Aliases](#).
8. Click on **OK** to complete the operation. The report opens and closes. The address is listed, its **Type** and **Status** depend on your configuration. For more details, refer to the section [Understanding the IP Address Type and Status](#).

Configuring and Managing IP Address Aliases

One IP address can have one unique FQDN that can be registered in the DNS. If additional names are necessary, you can register them as IP aliases.

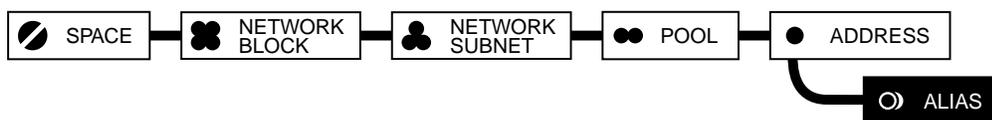


Figure 19.2. The alias in the IPAM hierarchy

There is no limitation for the number of IP aliases that SOLIDserver can manage. If you configured the update of the IP addresses in the DNS, the alias is usually a CNAME record created within the chosen domain that can therefore resolve the IP address name in your DNS servers.

The alias creation can be done from a free IP address or from a used IP address. Your alias can be named the way you want, its full name concatenates the name of one your existing domains to associate it with one your zones. Technically, the IP address alias can create either an A, AAAA or CNAME record in the DNS.

The aliases configuration can be used to point a record toward an IP address within one zone or toward an IP address saved in different zones. Within the same zone, the IP address alias is a CNAME record that follows the DNS standard use and points to an A/AAAA record. Among two different zones, the name is crucial: the IP address *shortname.domain1* creates an A record of the zone *domain1* and a CNAME record in the zone *domain2* with the value *shortname.domain2*. That way, your alias name links to two of your zones.

The most commonly used aliases create CNAME records in the DNS but, depending on the DNS configuration you want to set, you might need to create A records.

To let users follow the procedures below you need to configure the IPAM to DNS advanced properties to make sure the alias creation from the IP addresses creates records in the DNS. At network level, or higher, you need to choose a DNS server, set a *Domains list* and tick the box *Update DNS*. You can also set a default domain. For more details, refer to the chapter [Managing Advanced Properties](#). That way, the addresses you configure with aliases all inherit the properties and the records are successfully created in the DNS. Administrators can choose to display *All* advanced properties to set these options and successfully create records in the DNS.

To configure an alias on a free IP address

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Click on the **Name** of the available IP address of your choice. The pop-up window **This address is free, do you want to assign it?** opens.
4. Click on **OK**. The wizard opens.
5. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Shortname**, name the IP address.
7. In the drop-down list **Domain**, select a zone. The list of available zones depends on your configuration at network level.
8. In the drop-down list **Advanced properties**, select **All** and make sure that the property **DNS server** is configured with the value *All*.
9. Click on **NEXT**. The page **Aliases configuration** opens.
10. In the field **Name**, name your alias. Its name must be different from the IP address name, especially if they share the same domain.
11. In the drop-down list **Domain**, select an existing domain or *None*. The alias full name is displayed in the field **Alias** following the format: *<name>.<domain>* .
12. In the drop-down list **Type**, select *CNAME*, *A* or *AAAA*. By default, *CNAME* is selected.
13. Click on **ADD** to move your alias to the **Aliases list**. Repeat these actions for as many aliases as you need. In the list, each alias is listed as follows: *(<record-type> <full-alias-name>* .
14. Click on **OK** to complete the operation. The report opens and closes. The address is listed, its **Type** and **Status** depend on your configuration. For more details, refer to the section [Understanding the IP Address Type and Status](#).

To see its aliases in the list, display the column **Aliases**. For more details, refer to the section [Customizing the List Layout](#).

To configure an alias on an IP address in use

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the IP address of your choice, click on **E**. The properties page opens.
4. In the panel **Aliases**, click on **EDIT**. The wizard opens.
5. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The page **Aliases configuration** opens.
7. In the field **Name**, name your alias. Its name must be different from the IP address name, especially if they share the same domain.
8. In the drop-down list **Domain**, select an existing domain or *None*. The alias full name is displayed in the field **Alias** following the format: *<name>.<domain>*.
9. In the drop-down list **Type**, select *CNAME*, *A* or *AAAA*. By default, *CNAME* is selected.
10. Click on **ADD** to move your alias to the **Aliases list**. Repeat these actions for as many aliases as you need. In the list, each alias is listed as follows: *(<record-type>) <full-alias-name>*.
11. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again, the panel **Aliases** lists all aliases.

To edit an IP address alias

1. Go to the properties page of the IP address of your choice. For more details, refer to the procedure [To display an IP address properties page](#).
2. In the panel **Main properties**, click on **EDIT**. The wizard opens.
3. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Click on **NEXT**. The page **Aliases configuration** opens.
5. In the field **Aliases list**, select the alias you want to edit. The alias details are displayed in each of the relevant fields.
6. Make the changes you need.
7. Click on **UPDATE**. The alias is edited and listed in the **Aliases list**.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again, the panel **Aliases** lists all aliases.

To remove an IP address alias

1. Go to the properties page of the IP address of your choice. For more details, refer to the procedure [To display an IP address properties page](#).
2. In the panel **Main properties**, click on . The wizard opens.
3. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on . The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Click on . The page **Aliases configuration** opens.
5. In the field **Aliases list**, select the alias you want to remove and click on . The alias is no longer listed.
6. Click on to complete the operation. The report opens and closes. The properties page is visible again, in the panel **Aliases** the alias is no longer listed.

Configuring Multiple A Records for an IP Address

You can create several A records for one IP address from the IPAM module. That way, one IP address can have several aliases in the DNS, this can be especially useful when configuring load balancing and round-robin. For more details, refer to the section [Load Balancing with Round-Robin](#).

We strongly recommend against configuring your DNS with one IP address associated with a set of A aliases. Indeed, a proper configuration of your DNS implies that a name zone is configured with a reverse zone which allows DNS clients to query your domain, through its name on the one hand and its IP address on the other. In this configuration, DNS best practices advise to create a PTR record in the reverse zone for each A record of the name zone to make sure the domain or sub-domain is accessible through its name and IP address. If your name zone contains several A records with the same value, your reverse zone should contain as many PTR records. These records would all be named after the same IP address (the value of the A records). In this case, the reverse zone would contain several PTR records with the same name pointing to different domains. Therefore querying this IP address to get the corresponding domain or sub-domain is impossible: the server cannot know which hostname to send when answering the DNS clients query. To make sure that a domain can be accessed through its name and IP address, there should be one PTR record in the reverse zone for each A record of the name zone. If you need to provide an alias, you should add a CNAME record pointing to the A record in the master zone. For more details, refer to the sections [Adding an A Record](#), [Adding a AAAA Record](#), [Adding a PTR Record](#) and [Adding a CNAME Record](#).

To let users follow the procedures below you need to configure the IPAM to DNS advanced properties to make sure the alias creation from the IP addresses creates records in the DNS. At network level, or higher, you need to choose a DNS server, set a *Domains list* and tick the box *Update DNS*. You can also set a default domain. For more details, refer to the chapter [Managing Advanced Properties](#). That way, the addresses you configure with aliases all inherit the properties and the records are successfully created in the DNS. Administrators can choose to display *All* advanced properties to set these options and successfully create records in the DNS.

Keep in mind that if your configuration is not properly set in the IPAM, the A and/or AAAA records are not created in the DNS and no error message is displayed in the DNS.

To create several A/AAAA records when assigning an IP address

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Click on the **Name** of the available IP address of your choice. The pop-up window **This address is free, do you want to assign it?** opens.
4. Click on **OK**. The wizard opens.
5. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Shortname**, name the IP address.
7. In the drop-down list **Domain**, select a zone. The list of available zones depends on your configuration at network level.
8. In the drop-down list **Advanced properties**, select **All** and make sure that the property **DNS server** is configured with the value *All*.
9. Click on **NEXT**. The page **Aliases configuration** opens.
10. In the field **Name**, name your alias. Its name must be different from the IP address name, especially if they share the same domain.
11. In the drop-down list **Domain**, select an existing domain or *None*. The alias full name is displayed in the field **Alias** following the format: *<name>.<domain>* .
12. In the drop-down list **Type**, select *A* or *AAAA*.
13. Click on **ADD** to move your alias to the **Aliases list**. Repeat these actions for as many records as you need. In the list, each one is listed as follows: *(A) <full-alias-name>* or *(AAAA) <full-alias-name>*.
14. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again, the panel **Aliases** lists all aliases.

To create several A/AAAA records for an assigned IP address

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the IP address of your choice, click on **⚙**. The properties page opens.
4. In the panel **Aliases**, click on **EDIT**. The wizard opens.
5. If you or your administrator created classes at the IP address level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The page **Aliases configuration** opens.

7. In the field **Name**, name your alias. Its name must be different from the IP address name, especially if they share the same domain.
8. In the drop-down list **Domain**, select an existing domain or *None*. The alias full name is displayed in the field **Alias** following the format: *<name>.<domain>* .
9. In the drop-down list **Type**, select *A* or *AAAA*.
10. Click on to move your alias to the **Aliases list**. Repeat these actions for as many records as you need. In the list, each one is listed as follows: *(A) <full-alias-name>* or *(AAAA) <full-alias-name>*.
11. Click on to complete the operation. The report opens and closes. The properties page is visible again, the panel **Aliases** lists all aliases.

To edit or remove A and AAAA record aliases, refer to the procedures [To edit an IP address alias](#) and [To remove an IP address alias](#).

Renaming IPv4 Addresses Massively

SOLIDserver provides a tool to rename IPv4 addresses massively. There are three ways of renaming the IP addresses, you can replace part of the name or append/prepend the name with characters.

To massively rename IPv4 addresses

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. Tick the IP address(es) you want to rename.
4. In the menu, select **Edit > Replace > IP address name**. The wizard **Replace name of IP addresses** opens.
5. In the field **Exact search**, select one of the following options:

Table 19.6. Available methods to rename IP addresses

Option	Description
Replace	Select this option to rename the IP address entirely or partially. The fields <i>Replace</i> and <i>Name</i> are displayed.
	<i>Replace</i> : type in the name or part of the name to be replaced.
	<i>Name</i> : type in the characters that replace the value specified in the field <i>Replace</i> .
Append	Select this option to concatenate characters at the end of the name of the selected IP address(es). The field <i>Name</i> is displayed.
	<i>Name</i> : type in the characters you want to append to the name.
Prepend	Select this option to concatenate characters at the beginning of the name of the IP address selected. The field <i>Name</i> is displayed.
	<i>Name</i> : type in the characters you want to prepend to the name.

6. Click on to complete the operation. The report opens and closes. The new IP addresses names are visible in the list.

Moving IP Addresses

SOLIDserver provides several ways of moving, or migrating, IPV4 addresses within your database. Migrating IP address can be useful when you have to relocate hosts.

Keep in mind that migrating an IP address edits its class parameters' inheritance and propagation configuration: the value of the parameters is kept but each property configuration is forced to Set/Propagate.

Moving IPv4 Addresses to another Network

You can massively move IPv4 addresses from one terminal network to the other, i.e. you can relocate hosts and move them to another network. This operation allocates the first available IP address in the specified terminal network and keeps their properties and aliases.

Keep in mind that migrating an IP address edits its class parameters' inheritance and propagation configuration: the value of the parameters is kept but each property configuration is forced to Set/Propagate.

To move an IPv4 address to another network

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to move to another network.
4. In the menu, select **✎ Edit > Migrate to another network**. The wizard **Migrate addresses to another network** opens.
5. In the drop-down list **Target space**, select a space.
6. In the field **New network IP**, type in the start address of the terminal network you want to move the address(es) to.
7. Click on **[OK]** to complete the operation. The report opens and closes. The page is visible again, you can filter it to check the new address assigned to your hosts.

Moving IP Addresses to another Space

You can massively move IP addresses from one space to the other and keep their properties and aliases. To successfully migrate IP addresses to another space, keep in mind that:

- You can migrate IP addresses if the target space contains a terminal network with:
 - A start address and prefix that allow it to receive all the selected IP addresses. If the network is too small, the migration is impossible.
 - Enough free addresses to assign the IP addresses you migrate.
- You cannot migrate an IP address if the target network already has the same IP address *in use*. The migration would be interrupted.
- You cannot migrate an IP address from or towards a pool in read-only.

Keep in mind that migrating an IP address edits its class parameters' inheritance and propagation configuration: the value of the parameters is kept but each property configuration is forced to Set/Propagate.

To move IPv4 addresses to another space

1. In the sidebar, go to **♣ IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to move to another space.
4. In the menu, select **✎ Edit > Migrate to another space**. The wizard **Migrate addresses to another space** opens.
5. In the drop-down list **Target space**, select the space of your choice.
6. Click on **[OK]** to complete the operation. The report opens and closes. The page is visible again, display the addresses of the target space to see your addresses listed.

To move IPv6 addresses to another space

1. In the sidebar, go to **♣ IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to move to another space.
4. In the menu, select **⚙ Tools > Expert > Migrate to another space**. The wizard **Migrate addresses to another space** opens.
5. In the drop-down list **Target space**, select the space of your choice.
6. Click on **[OK]** to complete the operation. The report opens and closes. The page is visible again, display the addresses of the target space to see your addresses listed.

Migrating the Properties of an IPv4 Address

You can migrate the properties and aliases of one IP address individually. This allows to relocate a host to a different space or network or within the same space.

The wizard provides the possibility to detect an IP address in a source space and change the IP address itself, while keeping its properties in the target space.

Keep in mind that migrating an IP address edits its class parameters' inheritance and propagation configuration: the value of the parameters is kept but each property configuration is forced to Set/Propagate.

To migrate the properties of an IPv4 address

1. In the sidebar, go to **♣ IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **⚙ Tools > Migrate a specific IP**. The wizard **Migrating a specific IP address** opens.
4. In the drop-down list **Source space**, select a space or *Auto-detection*. This option can be selected only if your IP address exists in only one space.
5. In the field **IP address to migrate**, type in the IP address you want to migrate.

6. In the drop-down list **Target space**, select a space or *Same as source*. This option can be selected only if you specify a different IP address in the field *New IP address*.
7. In the field **New IP address**, type in the IP address of your choice. It can be a different one or the same as the one specified in the field *IP address to migrate*.
8. Click on to complete the operation. The report opens and closes. The page is visible again, you can filter it to check the new address assigned to your hosts.

Pinging an IP Address

From the IPAM module, you can ping IP addresses to check if the host they are associated with is responding.

The report can display the following messages:

- **Notice Ping OK (IP address).**

The corresponding host was found and responded to the ping.

- **Error Ping Timeout (IP address).**

The corresponding host did not respond to the ping. It could mean a number of things, the host is not running, is on a different network, is configured not to respond to the ping utility, etc.

To ping one or several IP addresses

1. In the sidebar, go to  **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to ping.
4. In the menu, select  **Tools > Ping**. The wizard **Pinging IP addresses** opens.
5. Click on to perform the ping. The report opens and displays the results.
6. In the section **Export format**, you can click on , or to export the result in the corresponding format. Even if you do not download the report, it is available in the window *Notifications* next to the field *Global search*.
7. Click on to go back to the page.

Deleting an IP Address

At IP address level, *deleting* an address actually frees it. Even though it is no longer listed, you can assign it again by search or manually. Note that, as deleting an address releases it, it is impossible to delete free addresses.

To delete an IP address

1. In the sidebar, go to  **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.

- Click on **OK** to complete the operation. The report opens and closes. Selected addresses are no longer listed.

Restoring an IP Address

You can restore deleted IP addresses, i.e. undo an IP address deletion. From the page *All deleted IP addresses*, you can restore your IP addresses.

From this page you can also export the entries listed or even create an alert or a chart.



Figure 19.3. The button undo on the page All Addresses

- This button allows to access the page **All Deleted IP addresses** of the version displayed on the page All addresses.

The page contains three columns:

Table 19.7. Columns on the page All deleted IP addresses

Column	Description
Date	The date and time of the IP address deletion.
User	The name of the user who deleted the IP address. This name is underlined, if you click on it you open user properties page in the module Administration.
Description	The IP address details: the IP address itself, its name, its MAC address (if relevant), the terminal network it belongs to and finally the space it belongs to.

To undo an IP address deletion

- In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
- On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
- On the right-end side of the menu, click on **undo**. The page **All deleted IP addresses** opens.
- Tick the IP address(es) you want to restore.
- In the menu, select **Edit > Undo IP deletion** or **Undo IP deletion (v6)**. The wizard **Restoring IP addresses** or **Restoring IPv6 addresses** opens.
- Click on **OK** to complete the operation. The report opens and closes. The addresses are no longer listed on the page **All deleted IP addresses**.

Automating the DNS Records Update From the Page All addresses

The name of an IP address as well as its aliases can automatically be created in the DNS via the advanced properties.

- The record is created in the zone you configured when setting the properties at network level. For more details, refer to the chapter [Managing Advanced Properties](#).
- Once the network properties are set and applied, to automate the record creation, you must tick the box *Update DNS*.

The records created is named after the IP address and zone name. Its value contains at least the IP address. For more details regarding IP address to DNS advanced properties, refer to the chapter [Managing Advanced Properties](#).

For more details regarding the aliases configuration, refer to the section [Configuring and Managing IP Address Aliases](#).

Automating the DHCP Statics Reservation From the Page All addresses

When you add IP addresses, you can automatically reserve a static in the DHCP thanks to the advanced properties.

- The static is created in the DHCP cluster you configured when setting the properties at network level. For more details, refer to the chapter [Managing Advanced Properties](#).
- Once the network properties are set and applied, to automate the static creation, you must tick the box *Create DHCP static*.

The new static shares the same IP address, MAC address and name as the new IP address. For more details regarding IP address to DHCP advanced properties, refer to the chapter [Managing Advanced Properties](#).

Updating Device Manager From the Page All addresses

At IP address level, there are two ways of updating Device Manager. You can use the option that populates its database or you can configure advanced properties to edit Device Manager database when adding IP addresses.

Populating Device Manager

From the page *All addresses*, you can select assigned addresses to create devices and interfaces in Device Manager. This operation creates a device that contains a set of interfaces based on each IP address and MAC address in your database.

For more details, refer to the section [Automatically Adding Devices from the IPAM](#) in the chapter Managing devices.

Editing Device Manager from the page All Addresses

From the page *All addresses*, you can update Device Manager, when adding an IP address. A set of advanced properties allow to create devices, associate an IP address with an existing device and edit the topology links between devices.

For more details, refer to the section [Manually Adding Devices from the IPAM](#) in the chapter Managing devices.

Chapter 20. Setting Up a Transition From IPv4 to IPv6

You can set up a semi-automated way to transition from IPv4 to IPv6 when creating IPv4 objects. You can now link the IPv4 networks or addresses you create with existing IPv6 networks or addresses as long as they belong to the same space. That way the day you stop using IPv4, your addressing plan is already configured with IPv6.

The transition option is managed like the advanced properties: you must configure and activate it.

Transition Specificities

The transition options are configured at space or network level and inherited at lower levels where you can apply them.

- For block-type networks, if the transition options are configured:
 - The transition to IPv6 can be set when adding or editing IPv4 networks.
 - The transition can only be set with existing IPv6 block-type networks.
- For subnet-type networks, if the transition options are configured:
 - The transition to IPv6 requires the address an existing block-type network, you either specify it or inherit from a configured parent network.
 - The transition to IPv6 automatically creates the appropriate IPv6 subnet-type network within the specified block-type network when you add or edit an IPv4 subnet-type network.
 - The IPv6 subnet-type network created is named after its IPv4 counterpart.
- For IP addresses, if the transition options are configured:
 - The transition is only possible if it was set at network level.
 - You can choose the IPv6 address creation behavior.
 - Adding or editing an IPv4 address creates the corresponding IPv6. The IPv6 address is named after the IPv4 address, has the same MAC address, device and class parameters.

If you edit an IPv4 subnet-type network already configured with a VLAN to set the transition to IPv6, the IPv6 corresponding network inherits the IPAM/VLAN interaction settings: both networks then belong to the VLAN.

Limitations

- The transition can only be set within one space: you cannot create IPv4 subnet-type networks in one space and expect to link them with IPv6 subnet-type networks in another space.
- If you set the transition parameters on an existing organization, they are not inherited and have to be set applied one object at a time.
- For block-type networks, the transition can only be configured and activated with existing IPv6 block-type networks. The transition options do not create block-type networks in IPv6 but only link IPv4 blocks with existing IPv6 blocks.
- The advanced properties set in IPv4 are not inherited by the corresponding IPv6 objects.

- At pool level, the transition options are not available.
- If you create an object in IPv4 and their IPv6 counterpart overlaps existing objects, only the IPv4 object is created.
- Deleting an IPv4 object linked to an IPv6 object does not delete the corresponding IPv6 object.

Configuring the IPv4 to IPv6 Transition

The transition has to be configured through the wizard *Advanced properties customization* and then activated. The configuration has to be set at space, network and IP address level. For the IP addresses the configuration set the IPv6 address creation behavior that suits your needs.

Configuring the Transition at Space and Network Level

The transition configuration is the same at space and network level: you need to display the IPv4 to IPv6 transition fields to be able to set the transition on every object.

To configure IPv4 to IPv6 transition at space and network level

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. In the menu, select **Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
3. Tick the box **Display the IPv4 to IPv6 transition fields**.
4. Click on **OK** to complete the operation. The report opens and closes. The page is visible again. Your configuration is now available in the addition and edition wizards.
5. In the breadcrumb, click on **All networks**. The page All network opens.
6. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
7. Repeat the steps 3 to 6 to configure the transition at network level.

Once the transition is configured, you can apply it when adding or editing spaces and IPv4 networks.

Configuring the Transition at IP Address Level

At IP address level, the transition option offer three ways of creating the IPv6 addresses:

1. **Offset** allows to take into account the position of the address you are assigning within the IPv4 terminal network and reuse it when assigning the corresponding IPv6 address. The 100th address of the IPv4 subnet-type network creates the 100th address of the related IPv6 subnet-type network.
2. **Injection** allows to convert in hexadecimal the IPv4 address you are assigning and use the whole address in its hexadecimal form as the last two bytes of the corresponding IPv6 address.
3. **First IP address available** allows to assign the first available IP address in the IPv6 terminal network when you assign an address in the IPv4 terminal network.

To configure IPv4 to IPv6 transition at IP address level

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.

2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, select **⋮ > Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
4. In the drop-down list **IPv4 addresses transition to IPv6 policy**, select *Offset*, *Injection* or *First IP address available*. By default, *Offset* is selected.
5. Click on **[OK]** to complete the operation. The report opens and closes. The page is visible again. Your configuration is now available in the addition and edition wizards.

Activating the IPv4 to IPv6 Transition

Once the option has been configured, you can activate it at every level of the IPAM hierarchy, except at pool level.

In the procedures below, the procedures use the inheritance from space to block-type network, block-type network to terminal network and terminal network to address. And create them one after the other. But each procedure can be used when editing the objects as long as you configured the option in the menu **⋮ > Extra options**.

The transition configuration details are available on the objects properties page, in IPv4 and IPv6.

Activating the Transition at Space Level

At space level, you can configure and apply the transition options. This does not create any IPv6 object but sets the existing IPv6 block-type networks of your choice for the transition. The options you set at this level are inherited by the IPv4 blocks you create in your space.

To activate the IPv4 to IPv6 transition when editing a space

1. In the sidebar, go to **♣ IPAM > Spaces**. The page **All spaces** opens.
2. Right-click over the **Name** of a space. The wizard **Edit a space** opens.
3. In the list **VLSM parent space**, select *None*.
4. Click on **[NEXT]**. The next page opens.
5. If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. Click on **[NEXT]**. The last page of the wizard opens.
6. Tick the box **Activate the IPv4 to IPv6 transition**. The field *IPv6 network (block)* appears.
7. In the field **IPv6 network (block)**, specify the beginning of the address of an existing block-type network in the space. The value set in this field must not exceed the first 2 bytes of the existing IPv6 block-type network. You cannot use the semi-colon (;) twice. That network you specify becomes the container of the IPv6 terminal subnet-type networks and addresses you create.
8. Click on **[OK]** to complete the operation. The report opens and closes. The page **All spaces** is visible again. Your configuration is available on the space properties page in the panel **Advanced properties**.

Once the option is configured, the IPv4 networks and addresses you create within your block-type network inherit this option. You can untick the box **Activate the IPv4 to IPv6 transition** if you do not want to set a transition for some of your objects.

Activating the Transition at Network Level

At network level, the transition settings applied at space level are inherited automatically. You need to create at least a block-type network and then subnet-type network.

You can untick the box *Activate the IPv4 to IPv6 transition* if you do not want to set the transition for a particular network.

When creating a block-type network, the settings are displayed and taken into account automatically. To set the transition for existing IPv4 networks, you need to edit them and tick the box.

To activate the IPv4 to IPv6 transition when creating a block-type network

1. In the sidebar, go to **♣ IPAM > Spaces**. The page **All spaces** opens.
2. Click on the name of the space you applied the transition options to. The page **All networks** of the space opens.
3. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
4. In the menu, click on **+ Add**. The wizard opens.
5. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **[NEXT]**. The page **Add an IPv4 network** opens.
6. In the field **Network Name**, name the network.
7. In the field **Address**, type in the start address.
8. In the drop-down list **Netmask** select a netmask. The netmask value automatically edits the *Prefix*.
9. In the section **IPAM properties**, the box **Activate the IPv4 to IPv6 transition** is ticked. The parameter is *Inherit* and set to *Propagate*. If you want to untick the box, you must *Set* the Inheritance property and then untick the box.
10. The field **IPv6 network (block)** displays the value set at space level. The parameter is *Inherit* and set to *Propagate*. If you want to edit it, you must *Set* the Inheritance property to be able to edit the value in the field, the bytes entered must correspond to an existing IPv6 block-type network.
11. Click on **[OK]** to complete the operation. The report opens and closes. The network is listed. The configuration is displayed on the properties page of the IPv4 and the IPv6 networks, in the panel **Advanced properties**.

Once the transition is configured on a block-type network, it is inherited by the subnet-type networks in contains.

When creating a subnet-type network, whether terminal and not, the settings are displayed and taken into account automatically. To set the transition for existing IPv4 subnet-type networks, you need to edit them and tick the box.

To activate the IPv4 to IPv6 transition when creating a subnet-type network

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.

3. Click on the name of the block-type network you applied the transition options to. The page **All network** of that network opens.
4. In the menu, click on **+ Add**. The wizard opens.
5. In the drop-down list **Network type**, select *Subnet*. Click on **NEXT**. The next page opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Add an IPv4 network** opens.
7. In the field **Network Name**, name the network.
8. In the field **Address**, type in the start address. By default, the start address of the block-type network you selected is displayed in the field.
9. Set the network size:
 - a. In the drop-down list **Netmask**, select a netmask. The netmask value automatically edits the *Prefix*.
 - b. In the drop-down list **Prefix**, select a value if you did not choose a netmask. The prefix value automatically edits the *Netmask*.

The network size configuration is visible in the field **Comment**.

10. In the section **IPAM properties**, the box **Activate the IPv4 to IPv6 transition** is ticked. The parameter is *Inherit* and set to *Propagate*. If you want to untick the box, you must *Set* the Inheritance property and then untick the box.
11. The field **IPv6 network (block-type network)** displays the value set at higher level. The parameter is *Inherit* and set to *Propagate*. If you want to edit it, you must *Set* the Inheritance property to be able to edit the value in the field, the bytes entered must correspond to an existing IPv6 block-type network in the space.
12. The field **IPv6 network (subnet)** displays the start address and prefix of an IPv6 subnet-type network and prefix that is created along with the IPv4 network you are creating.
13. Click on **OK** to complete the operation. The report opens and closes. The network is listed. In IPv6, the subnet-type network is created as well and shares the same name. The configuration is displayed on the properties page of the IPv4 and the IPv6 networks, in the panel **Advanced properties**.

Activating the Transition at IP Address Level

At address level, all the configurations set at higher levels are inherited.

To activate the IPv4 to IPv6 transition when assigning an IP address

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Click on the **Name** of a terminal subnet-type network you applied the transition options to. The page **All addresses** of the network opens.
4. Click on the available IP address of your choice. The pop-up window **This address is free, do you want to assign it?** opens. Click on **OK**. The wizard opens.
5. If you or the administrator created classes at the IP addresses level, in the list **IP address class**, select a class or *None*. Click on **NEXT**. The page **Add an IPv4 address** opens.

6. In the field **MAC address**, you can type in the MAC address of your choice. The IPv6 address is also associated to this MAC address.
7. In the field **Corresponding IPv6 address**, the IPv6 address is displayed in gray. This IP address depends on the transition you set in the wizard *Advanced properties customization*. For more details, refer to the section [Configuring the Transition at IP Address Level](#).
8. In the field **Shortname**, name the IP address. The IPv6 address is named the same. The **IP address name** field displays the shortname you typed in.
9. Click on **NEXT**. The last page of the wizard opens.
10. Click on **OK** to complete the operation. The report opens and closes. The address is listed. In IPv6, the address is created as well, it has the same name and MAC address. The configuration is displayed on the properties page of the IPv4 and IPv6 addresses, in the panel **Advanced properties**.

Chapter 21. Managing IPAM Templates

The IPAM provides a template that allows to create fully preconfigured IPAM structures, templates, to create the IPAM resources that suit your needs.

Using templates successfully implies:

1. Creating a template class for each relevant level of the IPAM hierarchy.
2. Creating the template resources that suit your needs.
3. Applying your template class when creating your resources. The template class applies the template configuration to your resource.

For instance, you can create a template for a block-type network containing 3 networks, each managing 3 pools, and associate it with a template class. Selecting this class when adding a new block-type network overwrites its name with its template name and automatically creates the related child objects with their template name, size and organization.

Note that the prefix set in the template class is not applied to the object you create with this class. However, all the template content is created in your resource, so make sure that the resource you create has a prefix small enough to contain all the objects set in the template.

Creating Template Classes in Class Studio

To create IPv4 templates, you need to create network classes and pool classes and set them both as *template class*.

You can create several template classes if you want to set up different organizations. You can use class directories to differentiate them. These classes are listed in the IPAM addition/edition wizards as follows: `<directory-name>/<class-name> [template]`.

To create an IPAM template class

1. **Create a network class**
 - a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
 - c. In the menu, click on **+ Add**. The wizard **Add a new class** opens.
 - d. In the field **Filename**, name the class.
 - e. In the field **Sub directory**, you can type in a directory name. If it does not exist it is created.
 - f. In the drop-down list **Module**, select **IPAM**.
 - g. In the drop-down list **Type**, select **Network**.
 - h. In the section **Enable class**, tick the box.
 - i. Click on **OK** to complete the operation. The report opens and closes. The class is now listed on the page among the *IPAM* module classes and marked *Enabled* in the column **Status**.

2. Create a pool class

- a. In the menu, click on **+** **Add**. The wizard **Add a new class** opens.
- b. In the field **Filename**, name the class.
- c. In the field **Sub directory**, you can type in a directory name. If it does not exist it is created.
- d. In the drop-down list **Module**, select **IPAM**.
- e. In the drop-down list **Type**, select **Pool**.
- f. In the section **Enable class**, tick the box.
- g. Click on **OK** to complete the operation. The report opens and closes. The class is now listed on the page among the *IPAM* module classes and marked *Enabled* in the column **Status**.

3. Enable the classes as template class

- a. In the list **Class Studio**, tick the classes you just created.
- b. In the menu, select **Tools > Use as template class**. The wizard **Add template** opens.
- c. Click on **OK** to complete the operation. The report opens and closes. The classes are marked *yes* in the column **Template**.

You can rename or edit a template class as long as you are not using it already. For more details, refer to the chapter [Configuring Classes](#).

Creating IPAM Templates

The template mode allows to create templates at every containing level of the IPAM module in IPv4. Before going further, keep in mind that:

- The template of an element created in Template mode that must be associated with a class to be used in Normal mode. If you created a whole hierarchy in template mode but only associated a terminal network with the appropriate class template, only this template is available in Normal mode.
- A template class can only be applied to one resource in template mode. You can use it on as many resources as you need in normal mode.
- A template can be configured with advanced properties and class parameters, they are applied to the resources you create in normal mode and inherited by the resources they contain.
 - The Inheritance property and Propagation property of the templates' parameters and properties can be *Set/Propagate*, *Set/Restrict* or *Inherit/Propagate*.
 - You cannot configure the Inheritance property and Propagation property to *Inherit/Restrict* in template mode.
- You cannot overlap addresses in Template mode even though you might only associate a few resources with a template class and apply them in normal mode.

At space level, you cannot apply template classes. Space templates are only the entry point of the templates hierarchy.

At network level:

- The start address set for a template is overwritten by the one you specify in normal mode: it is only used to define the network size.
- The prefix set for a template is overwritten by the one you specify in normal mode but the objects it contains are created with the configuration set in template mode. So make sure that the resource you create in normal mode is big enough to receive all the objects configured in template mode.
- A block-type network template, once associated with an enabled template class, allows to create a block-type network with the same name, properties and content as the template (networks, pools and assigned addresses).
- A subnet-type network template, once associated with an enabled template class or if the template network belongs to a block-type network associated with one, allows to create subnet-type network with the same name, properties and content as the template (pools and assigned addresses).
- The IPAM/VLAN interaction is limited: you cannot configure a subnet-type network template to *Create a VLAN*, you can only associate it with an existing VLAN.

At pool level:

- A pool template, once associated with an enabled template class or if the template pool belongs to a network associated with one, allows to create a pool named after the template, with the same properties and content (assigned addresses).
- The start address set for a template is overwritten by the one you specify in normal mode. If you configured the template pool with a template class, you must set the same size for the pool the
 - If you configured the template pool with a template class, creating a pool in normal mode with this class requires that you set it with the exact same size. Any other size returns an error.
 - If the template pool belongs to a network configured with a template class, creating the network automatically creates the pool. The start address of the template pool is then only used to define an offset for the first IP address of the pool and define its size.

At IP address level, you cannot apply templates. However, if you assign addresses belonging to a network template or pool template, they are assigned when you create the resources using the template.

Creating a Space Template

To be able to apply templates at network or pool level, you need to create a space template.

A space template cannot be associated with a template class, therefore you must configure the space with all the parameters and properties that you want to propagate at lower levels both in template and normal mode.

You can create several space templates if you want.

To create a space in template mode

Only users of the group *admin* can perform this operation.

1. Display the template mode.

- a. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
 - b. On the right-end side of the menu, click on **{ } Template Mode**. The page opens and a red message under the menu indicates that you are in template mode.
2. In the menu, click on **+ Add**. The wizard **Add a space** opens.
 3. If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. That class is not a template class, it does not allow to automate a space configuration in normal mode. Click on **NEXT**. The page **Add a space** opens.
 4. In the field **Space name**, name the space. The name you choose cannot be reused in template or normal mode afterward.
 5. In the field **Description**, you can type in a description.
 6. You can fill in the advanced properties fields following the table below.

Table 21.1. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

Keep in mind that these properties have to be manually set on the space you create in normal mode, the space template cannot be associated with a template class. Your configuration is, however, inherited by the resource you create in the space in template mode.

7. Click on **OK** to complete the operation. The report opens and closes. The space is listed.

Once your space template is created, you can create the objects it contains.

Keep in mind that in Template mode, you need to follow the IPAM hierarchy: the *Orphan Networks* and *Orphan Addresses* containers do not exist.

Creating a Network Template

Once you created a space template, you can create the network templates it contains.

Keep in mind that:

- You can only use templates for networks if you created a class for them, enabled it and set it as *template class*. For more details, refer to the section [Creating Template Classes in Class Studio](#).
- You cannot overlap block-type networks or subnet-type networks, even in template mode.
- You can create a block-type network in template mode that is not configured with a template class if you intend to only use templates for subnet-type networks and/or pools.
- You can create a subnet-type network in template mode that is not configured with a template class:
 - If you intend to only use templates for pools.

- If you intend to organize the content of a block-type network created in template mode but not configured with a template class.
- Any resource added in the block-type network becomes part of the template and is created when you add the network in normal mode. So you do not need to associate all lower level levels with a template class, but you can.
- Any resource added in the subnet-type network becomes part of the template and is created when you add the network in normal mode. So you do not need to associate all lower level levels with a template class, but you can.

To create a template for a block-type network

Only users of the group *admin* can perform this operation.

1. Display the template mode.
 - a. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
 - b. On the right-end side of the menu, click on **{ } Template Mode**. The page opens and a red message under the menu indicates that you are in template mode.
2. Click on the **Name** of the space template of your choice. The page **All networks** opens.
3. In the menu, click on **+ Add**. The wizard opens.
4. In the drop-down list **Network type**, select *Block*. Click on **NEXT**. The next page of the wizard opens.
5. In the list **Network class**, select the template class of your choice. Click on **NEXT**. The page **Add an IPv4 network** opens.

If you want to create a template at lower level, select *None*. Your configuration is inherited in template mode, but the content of your network cannot be created when you add a block-type network in normal mode.

6. In the field **Network name**, name the network.
7. In the field **Description**, you can type in a description.
8. In the field **Address**, type in the start address.
9. In the drop-down list **Netmask** or **Prefix**, select the value of your choice. The netmask you choose automatically edits the prefix and vice versa. The final size is displayed in the field **Comment**.
10. Depending on the administrator's configuration you may be able to fill in the advanced properties fields. They might be inherited from the space. For more details, refer to the section [Configuring IPAM Advanced Properties](#).
11. Click on **OK** to complete the operation. The report opens and closes. The network is listed.

Once you created a template for block-type networks, you can create a template for subnet-type networks. You can create it manually or using the option *By search*.

To create a template for a subnet-type network

Only users of the group *admin* can perform this operation.

1. Display the template mode.

- a. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
 - b. On the right-end side of the menu, click on **{ } Template Mode**. The page opens and a red message under the menu indicates that you are in template mode.
2. Click on the **Name** of the template for block-type network of your choice. The page **All networks** of the network opens.
 3. In the menu, click on **+ Add an IPv4 network (subnet) by search**. The wizard opens.
 4. In the list **Network class**, select the template class of your choice. Click on **[NEXT]**. The page **Network size** opens.

If you do want to create a template at this level, select *None*.

5. Select a **Size**, **Prefix** or **Netmask** for your network. Selecting one value automatically changes the other two. Click on **[NEXT]**. The page **Search result** opens.
6. In the list **Network address**, select a start address. Click on **[NEXT]**. The page **Add an IPv4 subnet** opens.
7. In the field **Network Name**, name the network.
8. In the field **Description** can display the value set on the parent network. The parameter is *Inherit* and set to *Propagate*. If you want to edit it, you must *Set* the Inheritance property to be able to edit the value in the field.
9. The fields **Address** and **Prefix** display the values set on the page *Network size*.
10. The box **Terminal network** is ticked by default to create a terminal network that automatically has a gateway. Depending on your administrator's display configuration:
 - a. The field **Gateway** can be displayed and you can edit it.
 - b. The field **Gateway** can be hidden but it is created anyway, based on the gateway offset calculation set in the wizard *Advanced properties customization* at network level.

You can untick the box if you want to create a non-terminal network that can contain other networks, in this case the network does not have a *Gateway*. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#).

11. Depending on the administrator's configuration you may be able to fill in the advanced properties fields. They might be inherited from the parent network. For more details, refer to the section [Configuring IPAM Advanced Properties](#) in the chapter [Managing Advanced Properties](#).
12. Click on **[OK]** to complete the operation. The report opens and closes. The network is listed.

If you do not need pools, you can go straight to the section [Applying a Template](#).

Do not hesitate to assign addresses within the terminal network templates, they are automatically assigned in normal mode when the template is used.

Creating a Pool Template

Once you created a template for subnet-type networks, you can create a pool template.

Keep in mind that:

- You can only use pool templates if you created a class for them, enabled it and set it as *template class*. For more details, refer to the section [Creating Template Classes in Class Studio](#).

- Pools cannot overlap each other, even in template mode.
- You can create a pool in template mode that is not configured with a template class if you intend to organize the content of a network created in template mode but not configured with a template class.
- Any IP address assigned in a pool becomes part of the template and is created when you add the pool or network that contains the pool.

To create a pool template

Only users of the group *admin* can perform this operation.

1. Display the template mode.
 - a. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
 - b. On the right-end side of the menu, click on **{ } Template Mode**. The page opens and a red message under the menu indicates that you are in template mode.
2. Create a pool template.
 - a. Click on the **Name** of the terminal network template of your choice. The page **All addresses** opens.
 - b. In the breadcrumb, click on **All pools**. The page opens.
 - c. In the menu, click on **+ Add**. The wizard opens.
 - d. In the list **IP pool class**, select the template class of your choice. Click on **[NEXT]**. The page **Add an IPv4 pool** opens.

If you do not want to set a template at this level, select *None*.
 - e. In the field **Pool name**, name your pool.
 - f. In the section **Pool read-only**, tick the box if you want all the addresses the pool contains to be reserved.
 - g. In the field **Start address**, specify the first address of the pool. By default, the first address of the parent network is displayed.
 - h. In the field **End address**, specify the last address of the pool. By default, the last address of the parent network is displayed. Editing this field automatically edits the field *Size* field, and vice versa.
 - i. In the field **Size**, type in the number of addresses you want in the pool. Setting a size edits the *End address*, if you do not specify anything, the size is automatically calculated according to the start and end addresses.
 - j. Depending on the administrator's configuration you may be able to fill in the advanced properties fields. They might be inherited from the parent network. For more details, refer to the section [Configuring IPAM Advanced Properties](#).
 - k. Click on **[OK]** to complete the operation. The report opens and closes. The pool is listed.

Once you created all the pool template(s) you need, you can apply them in normal mode.

Applying a Template

To apply the templates in *normal* mode, you need to select a template class when creating networks and pools. Before applying a template, keep the following information in mind:

The template mode limitations

1. Template classes do not exist for spaces.
 - If you create a network or pool template with a specific configuration of behaviors, you need to create a space in *normal* mode, with the exact same configuration as the space containing your objects in *template* mode.
 - In *template* mode, all the configuration parameters that you set for a network can only apply if the space is also configured with them.
2. In *normal* mode, you cannot name the object to which you apply a template class.

When you create an object using a template class, you must name it in the wizard but its name is automatically overwritten to match the object to which the class is associated in template mode.

For instance, if in *template* mode you created a network named *France* that you associated with the template class *france*, when you apply this class to a network you add in *normal* mode, you can name it *FR* in the addition wizard but when you commit your creation, the new network listed is named *France*.

3. In *normal* mode, you cannot rename an object created using a template.

The network templates specificities

1. The start address of a network/pool set in *template* mode indicates the size.

When you define a network/pool start address in the addition wizard, this address is actually used and the size of the template object is applied based on the start address you set.

2. When you are configuring a network using a template, type in the start address and make sure that the size automatically calculated is greater than the size of the template. In other words, make sure the automatic size matching the network can contain the template. The template class does the rest.

For instance, if you are using a subnet-type network template, type in the name and the start address. SOLIDserver proposes a size going from the start address you specified to the very last available in the block-type network. If this represents 512 addresses and your subnet-type network template sets up 128 addresses, do not modify anything as your subnet-type network template size can be contained in the subnet automatically calculated. Only the first 128 addresses you need are included in the newly created subnet-type network.

The pool template specificities

1. The pool template allows to provision subnet-type networks. All the assigned addresses are saved and recreated.
2. In normal mode when you apply a pool class template, the pool you create must be set with the same *Size* and *Pool read-only* than the template pool.

The procedure below illustrates the use of template classes when adding a block-type template. You can adapt this procedure to any level.

To apply a template when creating a block-type network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard opens.
4. In the drop-down list **Network type**, select *Block*. Click on **NEXT**. The next page of the wizard opens.
5. In the list **Choose a space**, select the space in which you want to add the network. Click on **NEXT**. The next page of the wizard opens.
6. In the list **Network class**, select the class template of your choice. Click on **NEXT**. The page **Add an IPv4 Network** opens.

The class name format is as follows: *your-template-name [template]* or *your-sub-directory/your-template-name [template]*.

7. In the field **Network Name**, name the network. This name is overwritten by the name of the template for block-type network that you selected.
8. In the field **Description**, you can type in a description.
9. In the field **Address**, type in the start address.
10. Do not edit the fields **Netmask** and **Prefix**. By default, no matter its content, the template automatically overwrites the default values to apply the size that you set in the template class.
11. Depending on the administrator's configuration you may be able to edit or complete the network configuration, the template's advanced properties and class parameters configuration are automatically loaded.
12. Click on **OK** to complete the operation. The report opens and closes. The network is listed, click on its **Name** to display its content: it contains all the objects set in template mode.

Chapter 22. Using VLSM to Manage Your IPAM Network

The Variable Length Subnet Masking (VLSM) is a technique that allows network administrators to break down the IP address organization on different levels of spaces, networks or pools both in IPv4 and IPv6.

From the space level you can use the IPAM hierarchy to model the organization of IP resources and increase its capacity. There are two ways of using VLSM within the GUI, both can be used to delegate user rights.

SOLIDserver uses specific icons to display VLSM hierarchies:

Table 22.1. VLSM icons on the pages All spaces and All networks

Icon	Description
●	The dot, located left of the space or network icon, indicates that the object belongs to another object. It specifies the level of the space or network in the VLSM hierarchy: one dot for level 1, two dots for level 2, three dots for level 3, and so forth depending on how deep your organization is.
●●	
●●●	
⌘	This icon indicates that the subnet-type network is non-terminal i.e. using VLSM. In a space-based VLSM organization it indicates that it is linked to a block-type network of the child space. In a network-based VLSM organizations, it indicates that the network contains other subnet-type networks.
⌘	This icon indicates that the block-type network is part of a space-based VLSM organization and belongs to a level 2 space, or lower. It shows that the block-type network is linked to a non-terminal subnet-type network in the parent space, they both share the same name and size.

Choosing the Method That Suits Your Needs

SOLIDserver provides **two different implementation techniques** for VLSM:

- A [Space-Based VLSM Implementation](#).
- A [Network-Based VLSM Implementation](#).

In both cases, organizing your IPAM network using VLSM provides a way to delegate user rights as it allows to limit what users belonging to certain groups can display and manage.

You can set up an organization that uses both methods but there are some requirements to meet. For more details, refer to the section [Properly Using Both Methods Simultaneously](#).

Space-Based VLSM Implementation

Your IPAM network can organize spaces: you can create as many spaces as you need and connect them with each other to set up your organization. The VLSM hierarchy can be as deep as you need.

VLSM introduces a parent to child dependency relationship between two spaces. A child space is then attached and related to a parent space, they are *affiliated*.

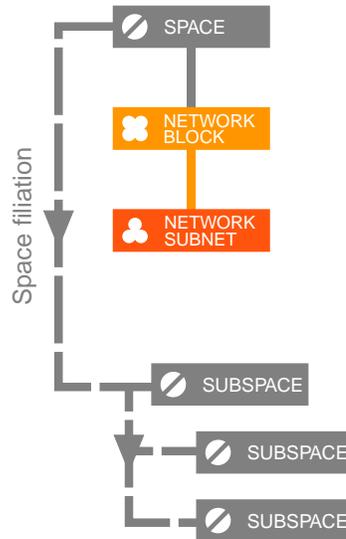


Figure 22.1. The space filiation in a space-based organization

The resources contained in the parent space can then be allotted to one of its child spaces. When a subnet-type network is created in a parent space, it may then be allotted to a child space: non-terminal subnet-type networks of parent space become block-type networks in the child space. This block-type network may then be divided into several subnet-type networks to be allotted as block-type networks in "grandchildren" spaces, and so on.

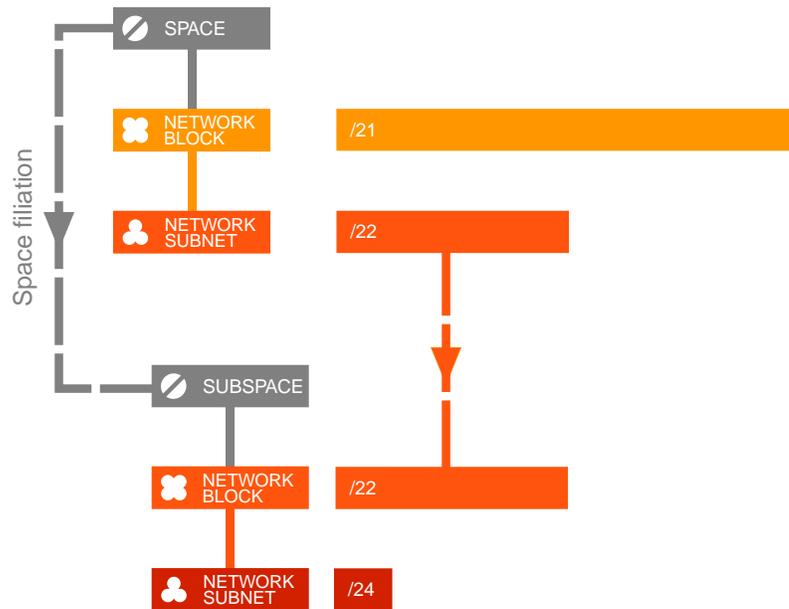
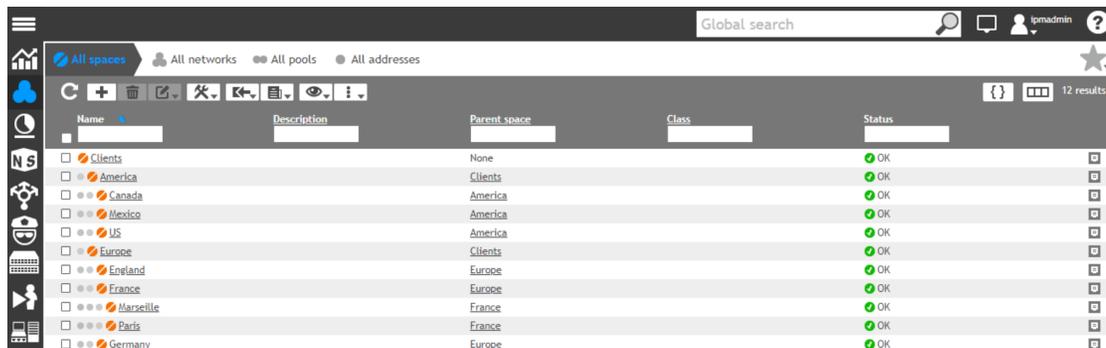


Figure 22.2. A delegation of networks among affiliated spaces

As spaces can be combined to map your organization, they can help network administrators to delegate the IP address management per layer of space. For instance, large block-type networks can be defined as root entries at the top level of the space hierarchy. These networks can be divided into several non-terminal subnet-type networks to be allotted to subspaces: each non-terminal subnet-type network becomes a block-type network in the child space. Within these subspaces, the block-type networks are divided into subnet-type networks matching the size of

your choice to register a network device, manage a specific set of IP addresses within your company...

This hierarchy makes it possible to obtain a coherent space unit where the resource administration is governed by the dependent relationships created between these spaces. The consistency check of resources and their uniformity are made between all affiliated spaces.



The screenshot shows the 'All spaces' page in an IPAM interface. The page displays a table of spaces with columns for Name, Description, Parent space, Class, and Status. The spaces are organized hierarchically, starting with 'Clients' at the top, followed by continents (America, Europe) and then specific countries (Canada, Mexico, US, England, France, Marseille, Paris, Germany). Each space has a status of 'OK'.

Name	Description	Parent space	Class	Status
clients		None		OK
America		Clients		OK
Canada		America		OK
Mexico		America		OK
US		America		OK
Europe		Clients		OK
England		Europe		OK
France		Europe		OK
Marseille		France		OK
Paris		France		OK
Germany		Europe		OK

Figure 22.3. A space-based VLSM organization on the page All spaces

In the example above, the clients IP database is organized based on geography: each country has a separate space affiliated to the continent it belongs to. These spaces were created prior to creating the networks and pools in order to shape the rest of the IP addresses organization.

With this type of organization, the delegation of rights can be set per continent, per country or within a country.

For more details regarding the manual VLSM implementation, refer to the section [Setting Up a Space-Based VLSM Organization](#).

Network-Based VLSM Implementation

Your IPAM network can organize subnet-type networks delegation within one space.

Using VLSM to Manage Your IPAM Network

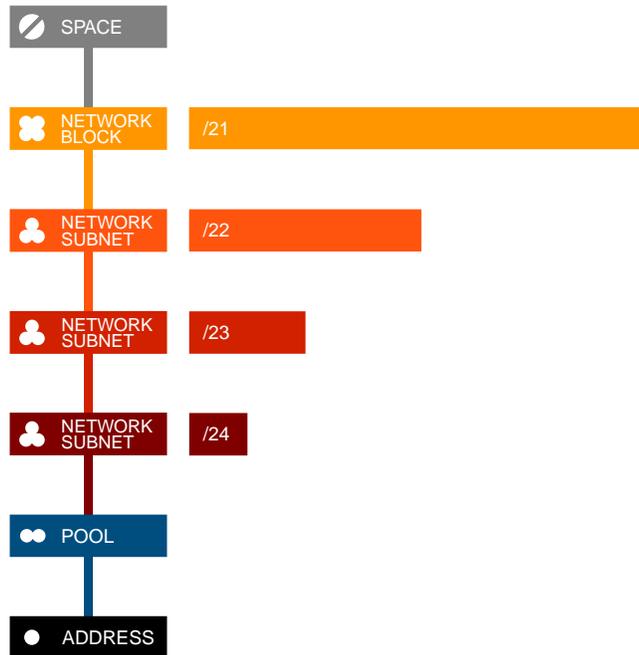


Figure 22.4. A network-based organization

Like the space-based implementation, it involves creating non-terminal subnet-type networks but this time it sets up several levels of hierarchy within one space. Therefore, the semi-automated VLSM allows you to distribute IP addresses on more than one level within a space without setting a space affiliation.

Address + prefix	Name	Size	Level	Space	Allocated (%)	Network Use (%)	Used IP (%)	Container	Status
2.0.0.0/8	internal	16777216	0	America	1.2%	0.4%	N/A	N/A	OK
2.1.0.0/16	managing-staff	65536	1	America	N/A	N/A	0%	internal	OK
2.2.0.0/15	staff-distribution	131072	1	America	0.8%	0.8%	N/A	internal	OK
2.2.0.0/24	new_york	256	2	America	N/A	N/A	0.4%	staff-distribution	OK
2.2.1.0/24	washington_dc	256	2	America	N/A	N/A	0.4%	staff-distribution	OK
2.2.2.0/23	san_francisco	512	2	America	N/A	N/A	0.2%	staff-distribution	OK

Figure 22.5. A network-based VLSM organization on the page All networks

In the example above, the *internal block* is divided into managing staff and all other staff members. The distribution of staff IP addresses is once again geographical. As the division is performed at the lower level, the delegation of rights to different administrators can be all the more precise with limited access to the database if necessary.

There is no limit to the number of non-terminal subnet-type network levels you can set. It all depends on their size and the size of the block-type network they belong to.

For more details regarding the network-based implementation, refer to the section [Setting Up a Network-Based VLSM Organization](#).

Properly Using Both Methods Simultaneously

To use both methods you must keep in mind that **VLSM implementation follows the IPAM hierarchy logic**: spaces contain block-type networks that contain subnet-type networks that

contain pools that contain IP addresses. Which is why you can only set both methods if you respect the following:

- You cannot set up a space-based organization using a space that already contains non-terminal subnet-type networks, meaning, a preexisting network-based organization. You can only set up first a space-based distribution and then a network-based delegation.
- The network-based implementation can only be set at the lowest of the space-based organization.

Keep in mind that, in a mixed organization, you can specify the inheritance and propagation properties depending on your needs. For more details, refer to the section [Editing a VLSM Block-type Network Class Parameters Inheritance](#).

Managing a Space-Based VLSM Organization

The space-based or manual VLSM implementation **must respect a specific order to be properly set**.

1. Create all the spaces one by one.
2. Affiliate all the spaces. Either when creating them or once created, as long as they are still empty.
3. Create all the block-type networks in the parent space(s).
4. Create the non-terminal subnet-type networks within these block-type networks. They become the block-type networks of the child space.

Once set up, you can move networks to edit the network links between a parent space and one of its children but you cannot set up a deeper organization once the spaces contain affiliated networks.

Note that you can also import space-based VLSM organizations, for more details refer to the section [Importing Spaces](#).

Setting Up a Space-Based VLSM Organization

Following the IPAM hierarchy, your space-based organization is done in the following order:

1. [Creating and Affiliating the Spaces](#) to set the delegation depth.
2. [Creating the Block-type Networks in the Parent Space](#) to define a range of IP addresses to delegate.
3. [Creating the Future Block-type Networks in the Child Space](#) that is to say creating non-terminal subnet-type networks in the parent space.

Keep in mind that **once the organization is implemented**:

- The network display is independent for each space level. From the page *All networks* of the top space you cannot display the networks of all space levels.
- Every non-terminal subnet-type network in a parent space becomes a block-type network in the child space.
- The content of a non-terminal network is common to both affiliated spaces: the terminal networks, pools and assigned IP addresses of a non-terminal subnet-type network are created both in the parent and child spaces.

- The subnet-type networks, pools and assigned IP addresses of a non-terminal subnet-type network are created both in the parent and child spaces.
- Any object you create in a parent space is created in the child and vice versa.
- Any object you delete in a parent space is deleted in the child and vice versa.
- Every terminal network in a space stays where it is created, even if it was created in a parent space.
- You cannot edit the space-based organization to make it deeper. However, some operations may relocate your networks. For more details, refer to the section [Editing a Space-Based VLSM Organization](#).

Creating and Affiliating the Spaces

The space affiliation is configured in the space addition or edition wizard.

Before creating your space affiliation, keep in mind that:

- You can have one parent space with several children. These children can also be parent to other spaces.
- The direct VLSM hierarchy affiliation between spaces updates simultaneously both spaces.
- The parameters and advanced properties configured at space level are inherited by the spaces, networks, pools and IP addresses they contain.

Therefore, the block-type networks of each space inherit the configuration, the subnet-type networks managed by the block-type network inherit as well, and so forth down to the IP addresses. You can edit any level of the hierarchy to set a particular parameter or property and propagate it to the lower level.

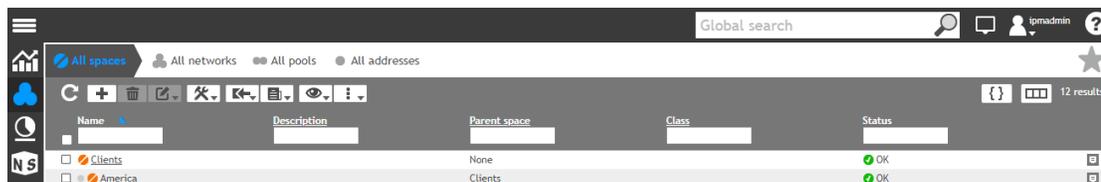


Figure 22.6. A space filiation where the space Clients contains the child space America

To set up the VLSM space-based organization

1. Creating the top level space

- In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
- In the menu, click on **+ Add**. The wizard **Add a space** opens.
- In the list **VLSM parent space**, select *None*. This first space is the top level space in your organization.
- Click on **NEXT**. The next page opens.
- If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.
- Fill the fields **Space name** and **Description** according to your needs.
- Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).

h. Click on **OK** to complete the operation. The space is listed.

2. Creating the affiliated child space

- a. In the menu, click on **+ Add**. The wizard **Add a space** opens.
- b. In the list **VLSM parent space**, select the top level space you just created.
- c. Click on **NEXT**. The next page opens.
- d. If you or your administrator created classes at space level, in the list **Space class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.
- e. Fill the fields **Space name** and **Description** according to your needs.
- f. Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).
- g. Click on **OK** to complete the operation. The level 1 space is listed.

You can repeat step 2 for as many spaces as you need. In the list **VLSM parent space**, you can select a space no matter its level in the hierarchy. The spaces display indicates each space level as detailed in the table [VLSM icons on the pages All spaces and All networks](#).

You can edit some aspects of the space-based organization even after you created networks, pools and addresses. For more details, refer to the section [Editing a Space-Based VLSM Organization](#).

Creating the Block-type Networks in the Parent Space

Once your space affiliation is set, you can create IPv4 or IPv6 block-type networks:

- The block-type networks must be created in the top level space of the affiliated spaces.

If you create them on a child space, the VLSM organization cannot be set, you would be creating a block-type network and not setting up a space-based organization.

- If you configured specific parameters or behaviors at space level, the block-type networks inherit them.

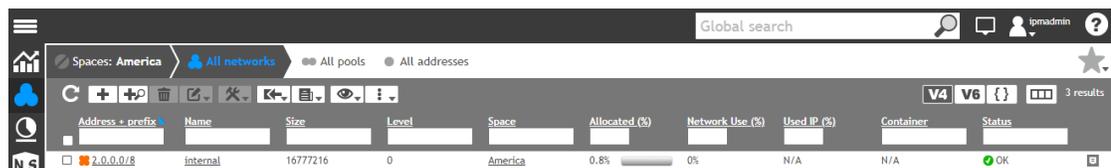


Figure 22.7. A block created in the space Clients

To add a block-type network manually

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard opens.

4. In the drop-down list **Network type**, select *Block*¹. Click on **NEXT**. The next page of the wizard opens.
5. In the list **Choose a space**, select the space in which you want to add the network. Click on **NEXT**. The next page of the wizard opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Add an IPv4 network** or **Add an IPv6 network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **Network Name**, name the network.
8. In the field **Description**, you can type in a description.
9. In the field **Address**, type in the start address.
10. If you are adding an IPv4 network:
 - a. In the drop-down list **Netmask** select a netmask. The netmask value automatically edits the *Prefix*.
 - b. In the drop-down list **Prefix**, select a value if you did not choose a netmask. The prefix value automatically edits the *Netmask*.

The network size configuration is visible in the field **Comment**.

11. If you are adding an IPv6 network, in the drop-down list **Prefix**, select a value between */16* and */64*. The values depend on the *Address* you specified.

If your administrator disabled the RFC 4291 compliance registry database entry, you can select a prefix between */16* and */128*. For more details, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

12. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 22.2. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

13. Click on **OK** to complete the operation. The report opens and closes. The network is listed.

To respect the levels of management you set, once you create a subnet-type non-terminal network in a top level space, if you click on its **Name** on the page *All networks* you are redirected to the content of the block-type network it created in the child space. This allows to navigate from one level to the other and create your terminal networks, pools and IP addresses.

¹If your group's permissions do not include the addition of both block-type and subnet-type networks, the page is automatically skipped.

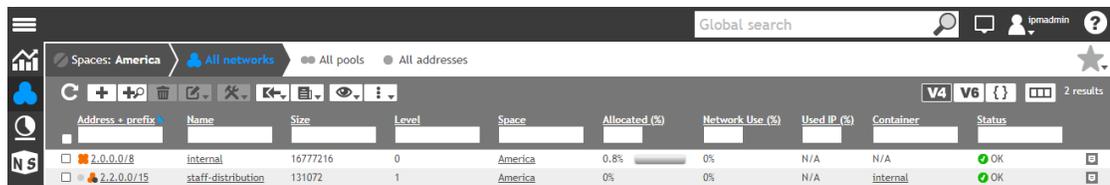
Once you created a block-type network in the parent space, you can create a non-terminal subnet-type network(s) it contains.

Creating the Future Block-type Networks in the Child Space

In the parent space, if you create non-terminal subnet-type networks in a block-type network they become the block-type networks of the child space of your choice. The terminal networks you create belong to the space, they are not delegated.

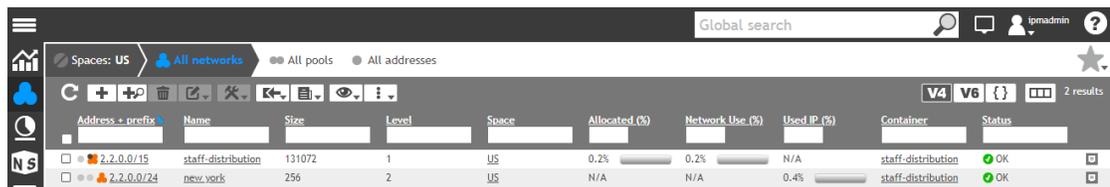
Keep in mind that, once you link block-type and subnet-type networks in a space affiliation, they update each other:

- If you create objects in the non-terminal subnet-type network, they are also created in the block-type network of the child space.
- If you create objects in the block-type network of the child space, they are also created in the subnet-type network of the parent space.



Address + prefix	Name	Size	Level	Space	Allocated (%)	Network Use (%)	Used IP (%)	Container	Status
2.0.0.0/8	Internal	16777216	0	America	0.8%	0%	N/A	N/A	OK
2.7.0.0/15	staff-distribution	131072	1	America	0%	0%	N/A	Internal	OK

Figure 22.8. A non-terminal network created in the block of the space America



Address + prefix	Name	Size	Level	Space	Allocated (%)	Network Use (%)	Used IP (%)	Container	Status
2.7.0.0/15	staff-distribution	131072	1	US	0.2%	0.2%	N/A	staff-distribution	OK
2.7.0.0/24	new_york	256	2	US	N/A	N/A	0.4%	staff-distribution	OK

Figure 22.9. In the child space US, the non-terminal network is a block-type network

In the procedures below, we create a non-terminal subnet-type network by search but you can also create it manually.

To create an IPv4 block-type network in a child space from a parent space

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. Click on the **Name** of the top level space. The page **All networks** opens.
3. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
4. Click on the **Name** of the block-type network of your choice to display its networks.
5. In the menu, click on **+ Add an IPv4 network (subnet) by search**. The wizard opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **[NEXT]**. The page **Network size** opens.
7. Select a **Size**, **Prefix** or **Netmask** for your network. Selecting one value automatically changes the other two. Click on **[NEXT]**. The page **Search result** opens.
8. In the list **Network address**, select a start address. Click on **[NEXT]**. The page **Add an IPv4 network** opens.

9. In the field **Network Name**, name the network.
10. In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).
11. The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
12. Untick the box **Terminal network**. The wizard refreshes and no longer includes the fields Gateway and pool related fields if they were displayed.
13. Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).
14. Click on **NEXT**. The page **VLSM space** opens.
15. In the list **VLSM space**, select the child space where the non-terminal subnet-type network becomes a block-type network.
16. Click on **OK** to complete the operation. The report opens and closes. The network is listed with the icon ♣. Depending on the organization depth, it is preceded by one or several ●. For more details, refer to the table [VLSM icons on the pages All spaces and All networks](#).

On the page **All networks** of the child space you chose, the non-terminal subnet-type network is listed as a block-type network.

To create an IPv6 block-type network in a child space from a parent space

1. In the sidebar, go to ♣ **IPAM** > **Spaces**. The page **All spaces** opens.
2. Click on the **Name** of the top level space. The page **All networks** opens.
3. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
4. Click on the **Name** of the block-type network of your choice to display its networks.
5. In the menu, click on ➕ **Add an IPv6 network (subnet) by search**. The wizard opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Network size** opens.
7. Untick the box **Terminal network**. The wizard refreshes and no longer includes the fields Gateway and pool related fields if they were displayed.
8. In the drop-down list **Network prefix**, select the value of your choice between *16 bits* and *64 bits*.
9. Click on **NEXT**. The page **Search result** opens.
10. In the list **Network address (v6)**, select a start address. Click on **NEXT**. The page **Add an IPv6 network** opens.
11. In the field **Network Name**, name the network.
12. The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
13. Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).
14. Click on **NEXT**. The page **VLSM space** opens.

15. In the list **VLSM space**, select the child space where the non-terminal subnet-type network becomes a block-type network.
16. Click on to complete the operation. The report opens and closes. The network is listed with the icon . Depending on the organization depth, it is preceded by one or several . For more details, refer to the table [VLSM icons on the pages All spaces and All networks](#).

On the page **All networks** of the child space you chose, the non-terminal subnet-type network is listed as a block-type network.

If you add another non-terminal subnet-type network in the parent space, a new block-type network is created in the child space.

Editing a Space-Based VLSM Organization

Once you set up a space-based organization, you can:

- Edit the VLSM parent of a space.
- Edit subnet-type networks to make them terminal or not in some specific cases.
- Edit VLSM block-type networks to inherit the value of a class parameter from a parent network or space.
- Delete subnet-type networks in some specific cases.

You cannot set up a deeper organization of spaces once they contain networks linked from one level to the other.

Editing the VLSM Parent of a Space

You can edit the VLSM parent of a space:

- Only if it does not already contain a block-type network that was created when a non-terminal subnet-type network was created at higher level.

For more details, refer to the section [Editing a Space](#).

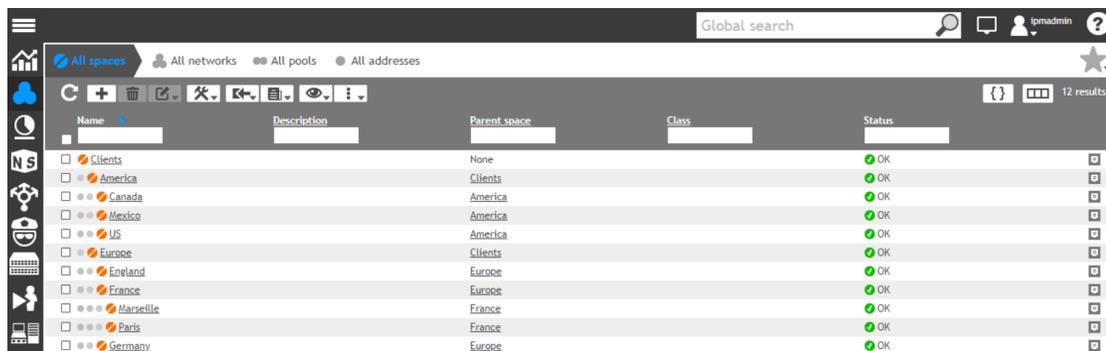
Editing a Subnet-type Network within a Space-Based Organization

Within a space-based organization, editing a subnet-type network terminal is limited:

- **Editing a terminal network to make it non-terminal**
 - Is only possible in the lowest level of the space organization, if the network contains pools. You must delete the pools and then edit the subnet-type network.
- **Editing a non-terminal network to make it terminal**
 - Is impossible if the block-type network created at lower level contains networks.
 - Is impossible if the non-terminal subnet-type network contains other networks, pools and/or IP addresses.
 - Is only possible if the block-type network created at lower level is empty.
- **Editing the target space of a subnet-type network created at lower level**
 - In an organization where a space manages several child spaces of same level, you can edit a non-terminal subnet-type network in the managing space to change its *VLSM space*. In the example below, this edition would be possible in the space *America*.

This edition moves the block-type network created in the former child space, along with all the networks, pools and addresses it contains, to the child space you specified.

- In an organization where a subnet-type network was created at several space levels - the spaces *Clients*, *Europe* and *France* in the example below - you can edit the VLSM space of the network and select the space *america*. That edition results in the creation of *Orphan containers* in the spaces *europa* and *france* and the creation of the block-type network in the space *america*. From there you can recreate the networks that suit your needs. With such an organization, this specific edition is only possible because the top space contained several spaces at the same level.



Name	Description	Parent space	Class	Status
<input type="checkbox"/> Clients		None		OK
<input type="checkbox"/> America		Clients		OK
<input type="checkbox"/> Canada		America		OK
<input type="checkbox"/> Mexico		America		OK
<input type="checkbox"/> US		America		OK
<input type="checkbox"/> Europe		Clients		OK
<input type="checkbox"/> England		Europe		OK
<input type="checkbox"/> France		Europe		OK
<input type="checkbox"/> Marseille		France		OK
<input type="checkbox"/> Paris		France		OK
<input type="checkbox"/> Germany		Europe		OK

Figure 22.10. A space-based VLSM organization on the page All spaces

Editing a VLSM Block-type Network Class Parameters Inheritance

Once you set a space-based VLSM organization, you can add and edit class parameters at the different levels of the hierarchy.

Both the parents of a VLSM block-type network, i.e. the space and the non-terminal subnet-type network, can have different values for the same class parameter. The VLSM block-type network inherits the parameter that has been configured before the other.

You can edit the class parameters inheritance source of a VLSM block-type network. The class parameter value can be inherited from:

- Space: the class parameter value is inherited from the lowest space in the organization, the space where the selected block-type network is located. You cannot choose a space located at higher level.
- Network: the class parameter value is inherited from the non-terminal subnet-type network it is linked with, one level up in the space hierarchy. You cannot inherit the value of any subnet-type network located higher in the hierarchy.

To edit a VLSM network class parameters inheritance

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. Tick the VLSM block-type network(s) of your choice.
3. In the menu, select **Tools > Expert > Edit class parameters inheritance source (VLSM block)**. The wizard **Edit class parameters inheritance source** wizard opens.
4. In the drop-down list **Parameter**, select the class parameter which value you want the selected network(s) to inherit.
5. In the drop-down list **Inherited from**, select the inheritance source:

- a. *Space* if you want to inherit the value from the space containing the block-type network you selected.
 - b. *Network* if you want to inherit the value from the network linked with the block-type network you selected, one level up.
6. Click on **ADD**. The wizard refreshes. The class parameter, along with its source and restriction option are listed in the **Properties list**.
 7. Repeat steps 5 to 7 for as many parameters as needed.

You can edit the content of the *Properties list*. Select an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.
 8. Click on **OK** to complete the operation. The report opens and closes.

Deleting Networks in a Space-Based Organization

As the IP address management in space-based VLSM organization relies on subnet-type and block-type networks distributed on different levels, deleting the networks has to follow certain rules.

- **You cannot delete a block-type network directly from the space it belongs to.**

Deleting a VLSM block-type network can only be done from higher level: when you delete the non-terminal subnet-type network that created it, the block-type network is deleted as well.

- **You can only delete a subnet-type network if the block-type network it created at lower level is empty.**

You cannot delete a subnet-type network in a parent space if the block-type network created in the child space contains networks

Deleting subnet-type networks automatically deletes the block-type network it created at lower level.

Unifying the VLSM Networks

In some cases after a migration, you might have a space-based VLSM organization where the parent and child spaces content do not match. The following two options allow to unify the content of your spaces especially the link between block-type and subnet-type network networks of two different levels.

Attaching a Network to its VLSM Parent

After a migration, the content of the parent and child spaces might differ and some objects might not be associated: you might have a block-type network in a child space that is not associated with a non-terminal subnet-type network.

Depending on your configuration, you might be able to create the missing non-terminal subnet-type network in the parent space using the option *Attach network to its VLSM parent*. This would allow to continue using VLSM to delegate rights and resources or create and delete objects in both spaces at once.

Keep in mind that you can only use this option if:

- The block-type network in the child space is not already associated with a non-terminal subnet-type network of the parent space.
- The parent space can receive the child space block-type network as a non-terminal subnet-type network:
 1. In the parent space, a block-type network can receive the non-terminal subnet-type network.
 2. There is no overlap.

To attach a network to its VLSM parent

1. In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
2. Click on the name of the child space of your choice. The page **All networks** open.
3. Tick the block-type network(s) of your choice.
4. In the menu, select . **Tools** > **Expert** > **Attach network to its VLSM parent**. The **Attach network to its VLSM parent** wizard opens.
5. Click on to complete the operation. The report opens and closes. The selected block-type network is now preceded by .

Aggregating VLSM Networks

After a migration, you can have affiliated spaces which IP addresses are not properly associated: you might have a missing non-terminal subnet-type network in a parent space even if the block-type network does exist in the child space.

In this case, you can use the option *Aggregate VLSM networks* from the page **All spaces** to create the missing non-terminal subnet-type networks in the parent space.

To aggregate VLSM networks

1. In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
2. In the menu, select . **Tools** > **Expert** > **Aggregate VLSM networks**. The **Aggregate VLSM networks** wizard opens.
3. In the drop-down list **Parent space**, select the parent space where the corresponding non-terminal subnet-type network is missing.
4. Click on to complete the operation. The report opens and closes. The missing non-terminal subnet-type network(s) is created on the page **All networks** list of the space you selected.

Moving IPv4 Addresses across the VLSM Hierarchy

If you reorganize your space-based hierarchy and plan on adding sub-spaces, you might need to move IPv4 terminal networks from a top level space to a lower level of the hierarchy. In that case, the IP addresses must be moved to the sub-space(s). This operation could be performed by moving IP addresses from one space to the other, as detailed in the section [Moving IP Addresses to another Space](#); however, you might have a lot of IPv4 terminal networks to spread on multiple sub-spaces, and it would take a long time to repeat the operation for each sub-space.

The option *Move addresses to VLSM network* allows to automate the migration of IPv4 addresses to the lowest terminal networks across the space hierarchy. It spreads the IP addresses in all the available terminal networks that can contain them. That is to say, a terminal network at the lowest level of the hierarchy which start address can receive the selected addresses.

To spread IPv4 addresses across the VLSM hierarchy

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Tick the IP address(es) you want to spread within the lower spaces of your VLSM organization.
4. In the menu, select **Tools > Expert > Move addresses to VLSM network**. The **Move VLSM IP addresses** wizard opens.
5. Click on **[OK]** to complete the operation. The report opens and closes. The IP addresses are listed, the space and networks they belong to have changed.

Managing a Network-Based VLSM Organization

The network-based VLSM organization organizes subnet-type networks imbrication within one space, it can be set within different block-type networks.

You can insert non-terminal subnet-type networks in between a child subnet-type network and its parent as detailed in the section [Reparenting subnet-type networks](#).

Note that you can also import network-based VLSM organizations, for more details refer to the section [Importing Networks](#).

Setting Up a Network-Based VLSM Organization

Setting up a network-based VLSM organization implies creating non-terminal subnet-type networks that manage terminal networks, pools and IP addresses on different levels. Note that:

- You can make the network-based VLSM organization go as deep as you need.
- You can use the network-based organization to delegate management, limit permissions or grant access to specific parts of your network. You can also use it to configure areas of your network differently, like in the image below.
- You cannot set up a network-based VLSM organization in a parent space. The network-based VLSM organization implies creating non-terminal subnet-type networks:
 - Either in an independent space,
 - Or in a child space that is at the lowest level of a space-based VLSM organization. In other words, in a child space that does not have any child space.

Indeed, creating non-terminal subnet-type networks in a parent space creates block-type networks in a child space, in which case, your VLSM organization is space-based.

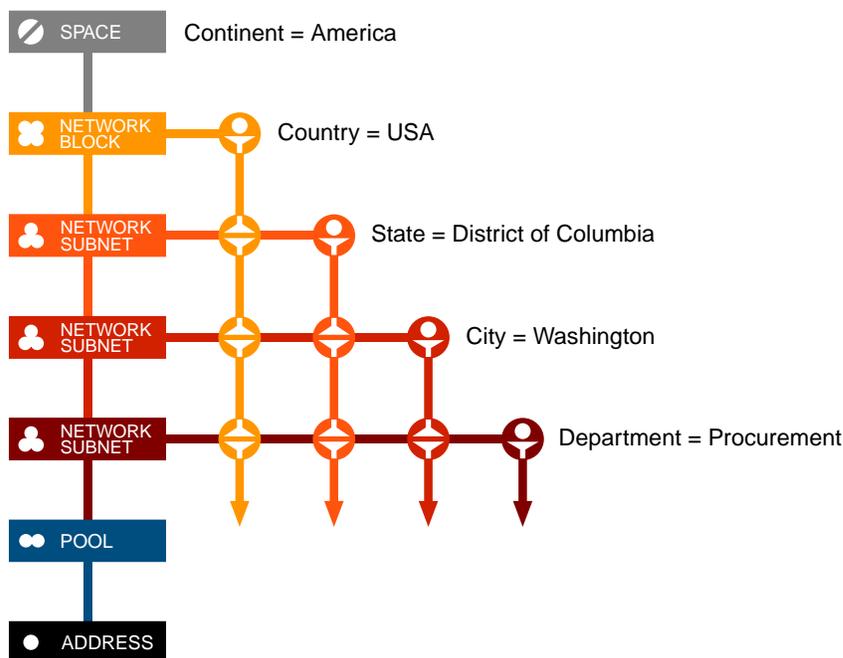


Figure 22.11. Example of a class parameter configuration based on the networks' organization

Contrary to the space-based organization, the network-based VLSM organization allows you to display the IPAM hierarchy at a glance, in one block-type network. As all the networks can be listed all together, there is no need to go through different spaces separately to list of non-terminal subnet-type networks and their content.

To create an IPv4 non-terminal subnet-type network

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. Click on the **Name** of the space of your choice. The page **All networks** of the space opens.
3. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
4. Click on the **Name** of the block-type network of your choice to display its networks.
5. In the menu, click on **+ Add by an IPv4 network (subnet) search**. The wizard opens.
6. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **NEXT**. The page **Network size** opens.
7. Select a **Size**, **Prefix** or **Netmask** for your network. Selecting one value automatically changes the other two. Click on **NEXT**. The page **Search result** opens.
8. In the list **Network address**, select a start address. Click on **NEXT**. The page **Add an IPv4 network** opens.
9. In the field **Network Name**, name the network.
10. The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
11. Untick the box **Terminal network**. The wizard refreshes and no longer includes the fields *Gateway* and *pool* related fields if they were displayed.

- Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).
- Click on to complete the operation. The report opens and closes. The non-terminal subnet-type network is listed and preceded by .

To create an IPv6 non-terminal subnet-type network

- In the sidebar, go to  **IPAM** > **Spaces**. The page **All spaces** opens.
- Click on the **Name** of the space of your choice. The page **All networks** of the space opens.
- On the right-end side of the menu, click on . The page refreshes and the button turns black.
- Click on the **Name** of the block-type network of your choice to display its networks.
- In the menu, click on  **Add an IPv6 network (subnet) by search**. The wizard opens.
- If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on . The page **Network size** opens.
- Untick the box **Terminal network**. The wizard refreshes and no longer includes the fields Gateway and pool related fields if they were displayed.
- In the drop-down list **Network prefix**, select the value of your choice between *16 bits* and *64 bits*.
- Click on . The page **Search result** opens.
- In the list **Network address (v6)**, select a start address. Click on . The page **Add an IPv6 network** opens.
- In the field **Network Name**, name the network.
- The fields **Address** and **Prefix** display the values set on the pages *Network size* and *Search result*.
- Using the drop-down list **Advanced properties**, you can configure advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).
- Click on to complete the operation. The report opens and closes. The non-terminal subnet-type network is listed and preceded by .

Once you created a non-terminal subnet-type network, it can contain as many non-terminal subnet-type networks as you need: it all depends on their size. And you can create a hierarchy as deep as you need.

Within the non-terminal subnet-type networks you can create terminal networks to manage pool and assign addresses. For more details, refer to the section [Adding Networks](#).

Reparenting Subnet-type Networks

The reparenting option allows you create a non-terminal subnet-type network and insert it between an already existing subnet-type network (terminal or non-terminal) and its parent. Therefore you can insert a new level in your network-based VLSM organization. Keep in mind that:

- The start or end address of the subnet network you are inserting must be different from the start and end address of the subnet-type network it reparents.

- This option can only be used in independent spaces or in spaces at the lowest level of a space-based VLSM organization. For more details, refer to the section [Properly Using Both Methods Simultaneously](#).

To reparent an IPv4 subnet-type network

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. Click on the **Name** of the space of your choice. The page **All networks** of the space opens.
3. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
4. In the menu, click on **+ Add**. The wizard opens.
5. In the drop-down list **Network type**, select *Subnet*². Click on **[NEXT]**. The next page of the wizard opens.
6. In the list **Choose a parent space**, select a non-terminal network among the ones listed under each space. The + sign left of the spaces' name opens the list of their networks.
7. Tick the box **Allow network reparenting**. Click on **[NEXT]**. The next page of the wizard opens.
8. If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on **[NEXT]**. The page **Add an IPv4 Network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. In the field **Network Name**, name the network.
10. In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).
11. In the field **Address**, type in the start address. By default, the start address of the block-type network you selected is displayed in the field. This address should be different from the existing subnet network that will be included in the network you are creating.
12. In the drop-down list **Netmask**, select a netmask. The netmask value automatically edits the *Prefix*.
13. In the drop-down list **Prefix**, select a value if you did not choose a netmask. The prefix value automatically edits the *Netmask*. The network size configuration is visible in the field
14. Untick the box **Terminal network**.
15. Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

Table 22.3. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer

²If your group's permissions do not include the addition of both block-type and subnet-type networks, the page is automatically skipped.

Field	Description
	to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The network is listed. The reparented network is now listed at a lower level in the VLSM hierarchy.

To reparent an IPv6 subnet-type network

- In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
- Click on the **Name** of the space of your choice. The page **All networks** of the space opens.
- On the right-end side of the menu, click on . The page refreshes and the button turns black.
- In the menu, click on **+ Add**. The wizard opens.
- In the drop-down list **Network type**, select *Subnet*³. Click on . The next page of the wizard opens.
- In the list **Choose a parent space**, select a non-terminal network among the ones listed under each space. The + sign left of the spaces' name opens the list of their networks.
- Tick the box **Allow network reparenting**. Click on . The next page of the wizard opens.
- If you or your administrator created classes at network level, in the list **Network class** select a class or *None*. Click on . The page **Add an IPv6 Network** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Network Name**, name the network.
- In the field **Description**, you can type in a description. The field may be in read-only if at a higher level, its Inheritance property is *Inherit*. If you want to specify a different description and/or restrict its propagation to lower levels, you must *Set* its Inheritance property and/or *Restrict* its Propagation property before being able to specify any value in the field. For more details, refer to the chapter [Inheritance and Propagation](#).
- In the field **Address**, type in the start address. By default, the start address of the block-type network you selected is displayed in the field. This address should be different from the existing subnet network that will be included in the network you are creating.
- In the drop-down list **Prefix**, select */64*, */127* or */128*.

If your administrator disabled the RFC 4291 compliance registry database entry, you can select a prefix between */16* and */128*. For more details, refer to the section [Enabling the Creation of IPv6 Terminal Networks with Non-Standard Prefixes](#).

- Untick the box **Terminal network**.
- Depending on the administrator's configuration, you may be able to fill in the advanced properties fields following the table below.

³If your group's permissions do not include the addition of both block-type and subnet-type networks, the page is automatically skipped.

Table 22.4. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the IPAM properties are detailed in the section Configuring IPAM Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The network is listed. The reparented network is now listed at a lower level in the VLSM hierarchy.

Editing a Network-Based VLSM Organization

You can edit a network-based organization by editing the networks themselves, or deleting non-terminal networks.

Editing a Subnet-type Network to Make it Terminal or Non-Terminal

To make a subnet-type network terminal or not, you must tick or untick the box **Terminal network** in the addition/edition wizard. For more details, refer to the section [Setting Up a Network-Based VLSM Organization](#) below.

Keep in mind that:

- **Making a terminal network non-terminal**

- You cannot edit a terminal network to make it non-terminal if it contains pools. You must first delete the pools.
- You can edit a terminal network to make it non-terminal even if it manages addresses. The assigned addresses it manages, including the gateway address, are moved to an *Orphan Addresses* container. This container is deleted once its content is deleted or the IP addresses are managed by another terminal network.

Note that if the network you edit belongs to a space that has child spaces, you can select a *VLSM space*. This action sets up a space-based delegation. In this case, a block-type network is created in the specified child space and it contains the same *Orphan Addresses* container.

- **Editing a non-terminal network to make it terminal**

- You cannot edit a non-terminal subnet-type network to make it terminal if it contains any network. It must be empty.

Deleting Subnet-type Networks in a Network-Based Organization

Deleting non-terminal subnet-type networks simply removes one level in the organization:

- In an organization with a non-terminal subnet-type network managing several terminal networks, these terminal networks and their content are managed directly by the block-type network.
- In a deep network-based organization, one level is removed. All the lower level networks it managed are moved up one level. If they inherited class parameters from the deleted container, for each class parameter:

- The Inheritance property is forced to *Inherit* or *Set* to match the configuration of the deleted parent network. The value and the source of the value remain the same.
- The Propagation property remains the same.

Using the VLSM Hierarchy to Delegate Management

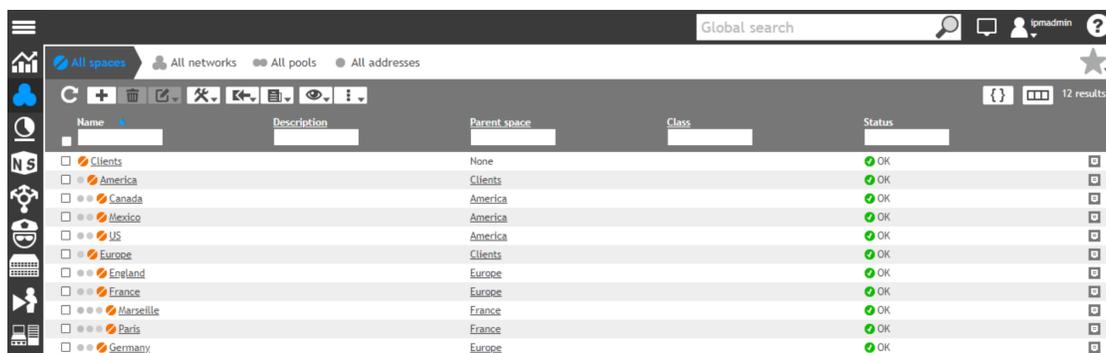
You can use the different levels of the VLSM hierarchy to organize, divide or limit the different users' rights in the IPAM module. Through the menu *Edit* menu you can make some spaces, networks or pools a resource for as many groups of users as you need. That way you can give them the possibility to add/delete/duplicate/move... the objects they contain.

If you make the different pieces of a space organization resources to specific groups, you can delegate the management one level at a time and whoever has access to the whole hierarchy can keep track of all the changes.

For more details regarding users, groups and delegation within SOLIDserver, refer to the part [Rights Management](#).

Delegating Management in a Space-Based Organization

In a space-based VLSM organization, the delegation could be done as follows.



Name	Description	Parent space	Class	Status
Clients		None		OK
America		Clients		OK
Canada		America		OK
Mexico		America		OK
US		America		OK
Europe		Clients		OK
England		Europe		OK
France		Europe		OK
Marseille		France		OK
Paris		France		OK
Germany		Europe		OK

Figure 22.12. An example of a space-based VLSM organization

In the example above, a way of using the space-based VLSM hierarchy could be to grant a group of users:

- Access and management permissions to the content of the space *usa*. That is to say, make all the block-type networks of the space *usa* a resource of the group and grant them the relevant permissions of these objects. That way, the users of the group:
 - Can see the space *usa*, as it is the container of the block-type networks they have in the page *Resources*.
 - Can display the content of the block-type networks they have among their resources. But cannot display the content of the subnet-type networks and pools of the block-type networks as they are not listed among their resources.
 - Can manage the block-type networks listed among their resources, as they were granted the relevant permissions.

Note that, if you grant the group access to the subnet-type networks and pools of the space *usa*, they can actually see the whole content of the space from the block-type networks down to the IP addresses and not only the subnet-type networks of the block-type networks. With

the relevant permissions, they can manage these objects as well. For more details, refer to the section [Assigning Resources to a Group](#).

- Grant access to the space *america*. That is to say, make the space *america* a resource of the group. That way, the users of the group:
 - Can display the content of the space *america* : all the spaces it contains, including the space *usa*. And therefore, see that all changes performed in the spaces *usa* and *america*.

Keep in mind that granting access to a resource does not grant users access to its container. Therefore, if you grant access to the group *usa*, users cannot see the content of *america*. In the same ways, if you grant access to *america*, users cannot see the content of *clients*.

If you grant access to these spaces to two different groups of users: one can perform specific operations and the other one to supervise these operations.

For more details regarding users, groups and delegation within SOLIDserver, refer to the part [Rights Management](#).

Delegating Management in a Network-Based Organization

In a network-based VLSM organization, the delegation could be done as follows.

Address - prefix	Name	Size	Level	Space	Allocated (%)	Network Use (%)	Used IP (%)	Container	Status
3.0.0.0/18	East Coast	16384	0	US	12.5%	4.7%	N/A	N/A	OK
3.0.0.0/22	New York	1024	1	US	25.0%	25.0%	N/A	East Coast	OK
3.0.0.0/24	New York City	256	2	US	100%	100%	N/A	New York	OK
3.0.0.0/25	Building 1	128	3	US	N/A	N/A	0.8%	New York City	OK
3.0.0.128/25	Building 2	128	3	US	N/A	N/A	0.8%	New York City	OK
3.0.4.0/22	Florida	1024	1	US	50.5%	50.5%	N/A	East Coast	OK
3.0.4.0/23	Miami	512	2	US	N/A	N/A	0.2%	Florida	OK
3.0.6.0/30	Orlando	4	2	US	N/A	N/A	50.0%	Florida	OK
3.0.6.4/32	Jacksonville	1	2	US	N/A	N/A	100%	Florida	OK
3.1.0.0/18	West Coast	16384	0	US	25.0%	18.8%	N/A	N/A	OK
3.1.0.0/20	California	4096	1	US	75.0%	75.0%	N/A	West Coast	OK
3.1.0.0/22	San Francisco	1024	2	US	N/A	N/A	0.1%	California	OK
3.1.4.0/22	Los Angeles	1024	2	US	N/A	N/A	0.1%	California	OK
3.1.8.0/22	San Diego	1024	2	US	N/A	N/A	0.1%	California	OK

Figure 22.13. An example of a network-based VLSM organization

In the example above, a way of using the VLSM hierarchy would be to grant a group of users:

- Access and management permissions to the all the subnet-type networks of the non-terminal subnet-type network *outside paris*. That is to say, grant them relevant permissions and make the non-terminal subnet-type network *outside paris* a resource of the group. That way, the users of the group:
 - Can display the content of the subnet-type networks they have among their resources.
 - Can edit the objects listed among their resources if they were granted the proper permissions and rights.

For more details, refer to the section [Assigning Resources to a Group](#).

- Grant access to the non-terminal subnet-type network *Field workers*. That is to say, make the non-terminal subnet-type network *Field workers* a resource of the group. That way, the users of the group:

- Can display the content of the non-terminal subnet-type network *Field workers* and see that it contains another subnet-type network.
- Can see that all changes performed on this range of addresses.

You can also only grant access to these subnet-type networks to two different groups of users. That way, you would set a group of users to manage specific operations and the other one to supervise these operations.

Note that if you grant access to the non-terminal subnet-type network *Field workers*, users are able to list the subnet-type networks it contains but not the content of the terminal network of the non-terminal subnet-type network.

For more details regarding users, groups and delegation within SOLIDserver, refer to the part [Rights Management](#).

Part VI. DHCP

The Dynamic Host Configuration Protocol (DHCP) is a network protocol which role is to automate the assignment of parameters to the clients connecting to the network, from a valid IP address to specific DHCP options.

It allows to set up specific connection behaviors for current devices and new devices on the network. The connection can be temporary, through **dynamic allocation**, or permanent, through **fixed reservation**. If you want to import an existing DHCP configurations, refer to the part [Imports and Exports](#).

The DHCP grants users access to the network following four steps:

1. **Discovery:** the DHCP client (host) broadcasts a DHCPDISCOVER packet on its physical subnet (usually 255.255.255.255) to discover the available DHCP servers;
2. **Offer:** all the available DHCP servers receiving the request respond to the host with a DHCPOFFER packet containing their own IP address and valid connection settings, dynamic or fixed;
3. **Request:** the host sends a DHCPREQUEST packet to inform all the DHCP servers that offered an IP of the acceptance. That packet includes the IP address of the DHCP server delivering access by the host, the other servers can return the offered IP address to their pool of available addresses;
4. **Acknowledge:** the selected server sends all the configuration data to the host in a DHCPACK packet.

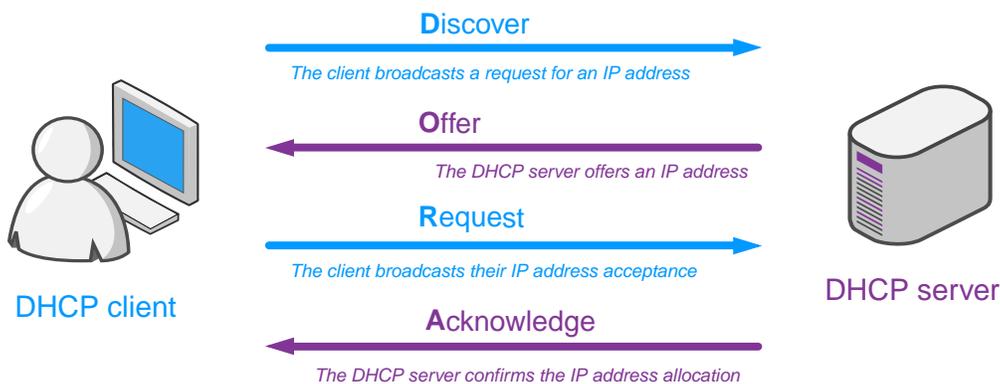


Figure 118. Representation of the DHCP IP address assignment process

The DHCP hierarchy can include up to 4 levels of organization. These levels depend on the connection behaviors you set, dynamic allocation or fixed reservation:

- **Servers:** the highest level of the hierarchy. It allows to set up fixed reservation and/or dynamic allocation. It can contain scopes, groups, ranges, leases and/or statics. For more details, refer to the chapter [Managing DHCP Servers](#). One or several servers can be managed via a smart architecture to ensure service availability and prevent data or configuration loss. For more details, refer to the chapters [Deploying DHCP Smart Architectures](#) and [Managing DHCP Smart Architectures](#).
 - **Groups:** an optional level of fixed reservation that belongs to a server and contains static reservations. It allows to apply specific DHCP options to the statics of your choice. For more details, refer to the section [Managing DHCP Groups](#) in the chapter [Managing Fixed Reservations](#).
 - **Scopes:** the second level of the hierarchy. It can contain ranges for dynamic allocation or directly statics for static allocation. For more details, refer to the chapter [Managing DHCP Scopes](#).
-

-
- **Ranges:** the third level of the hierarchy for dynamic allocation. They contain the IP addresses that are randomly allocated to hosts for a limited period of time, the leases. When the lease time expires, the address is returned to the range and can be reallocated. Ranges can be configured with Access Control Lists (ACLs) to restrict or authorize access to specific users. For more details, refer to the section in the [Managing DHCP Ranges](#) in the chapter [Managing Dynamic Addressing](#).
 - **Leases:** the lowest level of the hierarchy for dynamic allocation. Each lease is an IP address belonging to a range that is currently being allocated or was allocated to a host, you can track the leases history. For more details, refer to the section [Managing DHCP Leases](#) in the chapter [Managing Dynamic Addressing](#).
 - **Statics:** the lowest level of the hierarchy for fixed reservation. Static reservation ensures that a host always gets the same access details when they connect to the network. A static identifies a host using their MAC address and always provides them with an IP address and/or a set of DHCP options. For more details, refer to the section [Managing DHCP Statics](#) in the chapter [Managing Fixed Reservations](#).

The DHCP module also provides:

- **Failover channels.** The synchronization mechanism between two DHCP servers managing IPv4 addressing. It prevents allocation conflicts and depending on the configuration, it also allows disaster recovery. For more details, refer to the chapter [Managing Failover Channels](#).
- **DHCP options.** They can be set at any level of the hierarchy and are inherited at lower levels. For more details, refer to the chapter [Configuring DHCP Options](#).
- **IPv6 delegated prefixes.** They can be set on servers and linked to an existing DHCP shared network. For more details, refer to the chapter [Configuring DHCPv6 Prefix Delegation](#).
- **Monitoring and Reporting.** There are many ways for the reporting and monitoring of DHCP servers traffic and activity. For more details, refer to the chapter [Monitoring and Reporting DHCP Data](#).
- **Automation of IPAM and DNS resources creation.** The advanced properties allow to automate creations in the IPAM or DNS when you create DHCP resources. For more details, refer to the chapter [Managing Advanced Properties](#).

Note that from the module **Dashboards**, you can gather gadgets and charts on *DHCP dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 23. Deploying DHCP Smart Architectures

The DHCP can quickly become an essential piece of any network data organization. Once properly set up, it is usually hardly noticed, silently and faithfully performing its duties, day in and day out.

The DHCP clients' needs must be considered, including which DHCP options are supported by the client's operating system and which options and values need to be assigned. In large-scale DHCP implementations, the topology of the network becomes a very important factor. The network topology dictates where DHCP servers and/or relay agents must be placed. A final consideration is planning for fault tolerance.

To ensure that the DHCP service is available at all times and that you do not lose specific configurations if a DHCP server crashes, we strongly recommend that you manage physical servers through smart architectures. To understand the possible configurations of your service availability, refer to the section [Building a Highly Available Service With Smart Architectures](#).

Smart architecture pre-built DHCP configurations including backup and failover features with IPv4 addressing. Their deployment reduces the risk of misconfiguration.

There are several architectures and possible configurations to choose from both for DHCPv4 and DHCPv6. In DHCPv6, smart architectures simply provide a configuration backup.

For IPv4, four types smart architectures are available:

- *One-to-One*: in this DHCP configuration, two servers share the ranges of dynamic IP addresses.
- *One-to-Many*: this DHCP configuration is based on a central DHCP server with several peripheral DHCP servers as backup.
- *Split-Scope*: two DHCP servers are running in active/active mode and distribute the ranges management.
- *Single-Server*: this configuration manages one DHCP server. It provides a backup of the configuration that is pushed onto a new DHCP server if ever the original server crashed or stopped responding.

For IPv6, three types smart architectures are available:

- *Single-Server*: this configuration manages one DHCP server. It provides a backup of the configuration that is pushed onto a new DHCP server if ever the original server crashed or stopped responding.
- *Split-Scope*: two DHCP servers are running in active/active mode and distribute the ranges management.
- *Stateless*: this configuration provides a number of options to the servers managed through the architecture. The defined options, and not any other, are accessible to the DHCP clients. There is no limitation in the number of DHCP servers managed as this mode only provides options. Being based on DHCP options, stateless servers do not provide leases or include ranges or statics.

SOLIDserver supports a set of vendors detailed in the section [DHCP Vendors Compatible with Smart Architectures](#).

Implementing DHCP Smart Architectures

We strongly recommend that you manage every DHCP server with the smart architecture that suits your needs. Indeed, one of the main goals of this virtual management tool is to backup of your configuration. If the server(s) you are managing through the smart architecture were to crash, the architecture would save the configuration and allow you to push it on some new server(s) automatically. In addition, at any time you can change the type of smart architecture and the physical server(s) it manages.

DHCPv6 architectures have some particularities that differentiate them from the DHCPv4 architectures:

- The failover protocol is not available in IPV6. Thus, the page All failover channels in v6 is merely a list linking servers through the defined ports. For more details, refer to the chapter [Managing Failover Channels](#).
- IPv6 addressing is only possible from the EfficientIP servers.
- DHCPv6 servers operate on an appliance running in IPv4.
- In IPv6 there is no compatibility with the numerous vendors providing IP addressing

DHCPv4 One-to-One Smart Architecture

The DHCP One-to-One smart architecture allows you to quickly build a peer of two DHCP servers managing IPV4 addresses with a pre-built high availability mechanism. When you deploy a One-to-One smart architecture, you drastically reduce the DHCP service downtime if one of your DHCP servers is out of service.

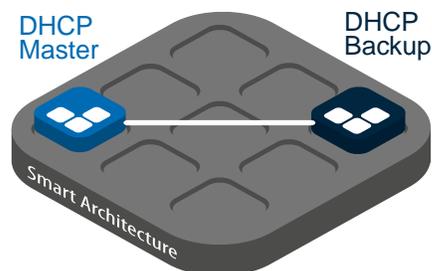


Figure 23.1. DHCPv4 One-to-One smart architecture

The One-to-One smart architecture allows two DHCP servers to share a range of common addresses. Should a server stop working, the second server would take over, depending on your failover configuration. For more details regarding failover, refer to the section [DHCP Failover Principles and Operational States](#) of the chapter [Managing Failover Channels](#).

DHCPv4 One-to-Many Smart Architecture

Functionally, the DHCP One-to-Many smart architecture is a replication of several One-to-One smart architectures, which is why it is only available for DHCP servers managing IPv4 addresses. The One-to-Many smart architecture is based on a set of DHCP servers that are all linked to only one Master DHCP server.

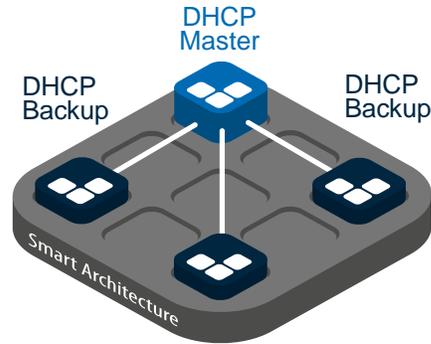


Figure 23.2. DHCPv4 One-to-Many smart architecture

This architecture is particularly relevant for organizations that have many sites and need to have a dedicated DHCP service per site. It looks like a star configuration, where N DHCP servers, no matter their location, share a failover channel with the central DHCP server of the smart architecture: it is a N+1 servers configuration.

DHCPv4 Split-Scope Smart Architecture

The Split-Scope smart architecture allows you to share ranges between two EfficientIP DHCP servers in an active/active configuration. You can define the proportion of IP addresses managed by each server. One server is set as a master and the other one as a backup. The main goal of this architecture is the availability of the services at all times thanks to the shared load.

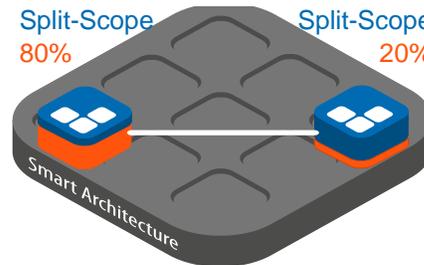


Figure 23.3. DHCPv4 Split-Scope smart architecture

There is no failover protocol between the two servers but being a smart architecture, the Split-Scope provides a backup of the configuration: if anything were to happen to any of the managed servers, installing them back to SOLIDserver would apply the smart architecture on both servers again.

DHCPv4 Single-Server Smart Architecture

The Single-Server architecture provides a backup of the management configuration of any of the available DHCP servers: EfficientIP DHCP, EfficientIP DHCP Package, Microsoft DHCP and Nominum DCS. Therefore, if it were to crash, you could install it again and let SOLIDserver push automatically the smart architecture configuration back onto your server.

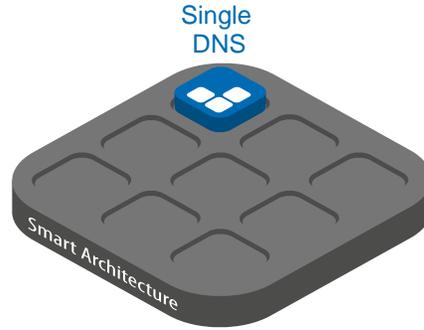


Figure 23.4. DHCPv4 Single-Server smart architecture

DHCPv6 Single-Server Smart Architecture

The DHCPv6 Single-Server architecture only provides a backup of the management configuration of an EfficientIP DHCP server. Therefore, if the physical server were to crash, you could install it again and let SOLIDserver push automatically the smart architecture configuration back onto your server.

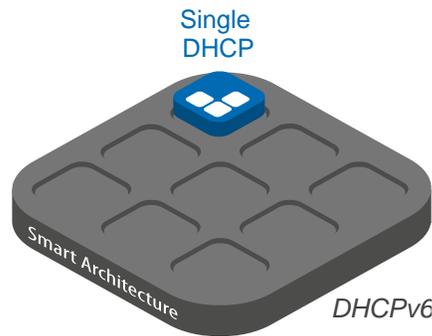


Figure 23.5. DHCPv6 Single-Server smart architecture

DHCPv6 Split-Scope Smart Architecture

The Split-Scope smart architecture allows you to share ranges between two EfficientIP DHCPv6 servers in an active/active configuration. You can define the proportion of IP addresses managed by each server. One server is set as a master and the other one as a backup. The main goal of this architecture is the availability of the services at all times thanks to the shared load.

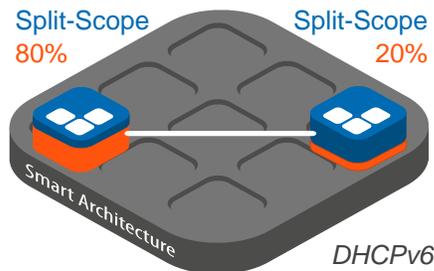


Figure 23.6. DHCPv6 Split-Scope smart architecture

There is no failover protocol between the two servers but being a smart architecture, the Split-Scope provides a backup of the configuration: if anything were to happen to any of the managed servers, installing them back to SOLIDserver would apply the smart architecture on both servers again.

DHCPv6 Stateless Smart Architecture

The Stateless smart architecture allows you to set up a number of options to the scopes of the servers you manage. The DHCP clients have access to a set of options that you define for the architecture. Which is why you can add as many servers as you need in this configuration.

There is no master or backup servers per se in this configuration. By default, they all are independent master servers sharing the same options configuration.

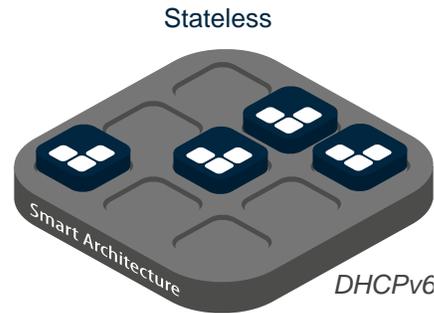


Figure 23.7. DHCPv6 Stateless smart architecture

Keep in mind that the Stateless smart architecture only has an impact on the options available to the DHCPv6 clients, therefore it is impossible to add ranges and static through this configuration. In the same way, no leases are provided or managed.

DHCP Vendors Compatible with Smart Architectures

SOLIDserver supports a set of [DHCP Servers Managing IPv4 Addressing](#). In [DHCP Servers Managing IPv6 Addressing](#), only EfficientIP servers are supported.

DHCP Servers Managing IPv4 Addressing

SOLIDserver appliances support both EfficientIP and other vendor DHCP servers, allowing you to configure and deploy IP services across your distributed network and synchronize data updates in real time from the GUI.

Table 23.1. DHCPv4 supported vendors

Vendor	Smart Architecture			
	One-to-One	One-to-Many	Split-Scope	Single-Server
EfficientIP DHCP	X	X	X	X
EfficientIP DHCP Package	X	X	X	X
Microsoft (all types)			X	X
Nominum DCS	X			X

SOLIDserver supports almost all features delivered by each vendor but does not add additional features at service level. Thus, limitations depend on each vendor. For instance, Microsoft Windows DHCP services do not provide failover, so you cannot configure it from the appliance.

You can manage any supported vendor on one page from the GUI. SOLIDserver is an abstraction layer that masks the specific processes of each DHCP vendor to network administrators. DHCP services are not managed one server at a time but as a global service. It is possible to simultaneously configure Microsoft Windows running DHCP servers and Linux running ISC DHCP servers, modify VoIP options on all DHCP servers or create transversal reports to get an immediate comprehensive understanding of network services configurations.

Each and every one of these servers can be managed by SOLIDserver smart architecture to ease the management configuration and provide a backup of the chosen configuration. For more details, refer to the chapter [Managing DHCP Smart Architectures](#).

DHCP Servers Managing IPv6 Addressing

With DHCPv6 addressing, the choice is more limited. For now, you can only manage EfficientIP DHCP servers but there are a number of architectures that allow you to manage either one or several EfficientIP DHCP servers at once.

For more details, refer to the chapter [Adding a DHCPv6 Smart Architecture](#).

Building a Highly Available Service With Smart Architectures

A way of maintaining DHCP service in case of a partial power loss or network outage is to set up two DHCP servers and enable them to both serve the same network. You can set up each server on different networks. In this case, if you lose connectivity or power on one network but not the other, the DHCP service continues.

Two active DHCP servers cannot share an IP address pool since they have no way of knowing with certainty which IP addresses are being distributed. Therefore, two active DHCP servers cannot perform dynamic DHCP which is why scope splitting is necessary to separate IP address ranges per server. This configuration, the Split-Scope, is available for both DHCPv4 and DHCPv6.

With a traditional active/passive pair of DHCP servers, if the active server fails, the network administrator is required to manually turn on the passive DHCP server so that it can take over until the initial active server is restored. DHCP High Availability with IP address scope splitting provides failover but with the risk of meeting downtime as addresses are leased to more than one client and have potential manual intervention to clean up the lease database. In order for two DHCP servers to provide DHCP services for the same network segments, the servers must coordinate their behavior. Each server must either know what the other is doing or be configured so that it can operate without knowing what the other is doing. In order for each server to know what the other is doing, the DHCP safe failover protocol can be implemented.

As the failover protocol is not available in DHCPv6, **the DHCP Safe failover protocol is only available for DHCPv4 servers.**

Chapter 24. Managing DHCP Smart Architectures

Once you chose the smart architecture(s) that suit your needs in the chapter [Deploying DHCP Smart Architectures](#), you can manage them following the sections below.

Browsing DHCP Smart Architectures

Smart architectures are managed from the page **All servers**, listed like physical servers and preceded by the icon . For more details, refer to the section [Browsing DHCP Servers](#).

Browsing the DHCP Smart Architectures Database

To display the list of DHCP smart architectures

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. To display the DHCP smart architectures and their content, click on  to display the physical server(s) managed by each smart architecture.

In the column **Name**, all the smart architectures are preceded by the icon . They are listed with the physical servers.

To display a DHCP smart architecture properties page

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. Filter the list if need be.
3. At the end of the line of the smart architecture of your choice, click on . The properties pages opens.

Understanding the Smart Architecture Statuses

The column **Status** provides information regarding the smart architectures' configuration.

Table 24.1. Smart architecture statuses

Column	Description
 <i>OK</i>	The smart architecture is operational.
 <i>Invalid settings</i>	The smart architecture does not contain any physical server, is missing one or several physical servers or is not configured properly (not enough failover channels configured, etc).
 <i>Locked synchronization</i>	The server configuration is not viable. For more details, refer to the section Handling the Status Locked Synchronization .

Moreover, the column **Sync** (i.e. synchronization) provides additional information regarding the exchanges between the smart architecture and the physical server(s).

Table 24.2. Smart architecture synchronization statuses

Column	Description
 <i>Synchronized</i>	The smart architecture has successfully synchronized the server(s) it manages.

Column	Description
 <i>Busy</i>	The smart architecture is synchronizing the server(s).
 <i>Locked synchronization</i>	The synchronizing failed as the server configuration is not viable: the smart architecture cannot send the configuration file to the physical server(s). For more details, refer to the section Handling the Status Locked Synchronization .

Adding a DHCPv4 Smart Architecture

A smart architecture can be configured without DHCP servers. It allows you to create the architecture that suits your needs before applying it to one or more DHCP servers. It also provides a backup of the management configuration of the server it manages. If your DHCP server crashes, you delete it and add a new one on which you apply the same architecture, SOLIDserver remembers the former server's configuration and apply it to the new one.

With DHCPv4, there are four different kinds of smart architectures: One-to-One, One-to-Many, Split-Scope and Single-Server. As for DHCPv6 smart architectures, SOLIDserver proposes the Single-Server, Split-Scope and Stateless architectures. In the procedures below, we are going to describe the configuration of the DHCP smart architectures with the DHCP servers they manage, but you can go through the configuration without adding any server and do it later. For more details, refer to the part [Adding a DHCP Server into a Smart Architecture](#).

Once the configuration is completed, the DHCP smart architecture appears on the page **All servers** as a real server.

As you can see, the column Type mentions the kind of smart architecture applied, the DHCP smart members column is marked N/A and for that reason, the server status is Invalid settings.

DHCPv4 One-to-One Smart Architecture

The [One-to-One Smart Architecture](#) allows you to set up a failover channel between two DHCP servers: one is set as master server and the other one as backup. Note that, if the master server crashes, you have to manually set up the partner-down mode to reclaim the available IP addresses, refer to the section [Operating in Partner-down State](#) in the chapter Managing Failover Channels. This architecture also provides a shared management of the leases that you can configure according to your needs.

To configure a DHCPv4 One-to-One smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 24.3. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.

Managing DHCP Smart Architectures

Parameter	Description
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DHCP smart architecture**, select **One-to-One**.

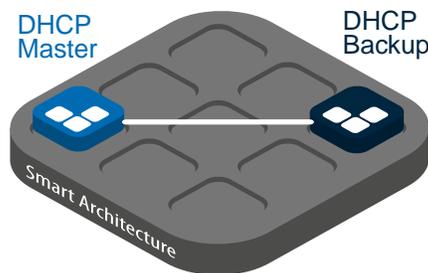


Figure 24.1. DHCPv4 One-to-One Smart Architecture

7. Click on **NEXT**. The next page of the wizard opens.
8. In the list **Available DHCP servers**, you can select one by one the two DHCP servers that you want to manage through the smart architecture.
9. Click on **+**. The selected server is moved to the list **Selected DHCP servers**. Repeat these actions for the second server.
10. Click on **NEXT**. The next page of the wizard opens.
11. In the drop-down list **Master DHCP server**, select the Master server in the smart architecture configuration.
12. Click on **NEXT**. The next page of the wizard opens.
13. This page allows you to configure the failover channel between the servers of the architecture. Fill in the fields according to the table below:

Table 24.4. DHCPv4 One-to-One failover parameters

Parameter	Description
Peering name	The default failover channel name, <i><failover-your.smart.server.name></i> , is displayed in the field. You can edit it.

Managing DHCP Smart Architectures

Parameter	Description
Failover port	Type in the number of the port on the master server dedicated to the failover. By default, the port 847 is used.
Failover peer port	Type in the number of the port on the backup server dedicated to the failover. By default, the port 647 is used.
Automatic switch to partner-down delay (in hours)	Type in this field the amount of time (in hours) after which a failover channel in Communications-interrupted state should automatically switch to Partner-down. The accepted values are between 4 and 65535. By default, the option is disabled and the field is set to 0.
Peer DHCP server	The DHCP backup server is automatically entered in this field.
Split leases	In this drop-down list, you can choose how to split the leases between the two servers: <i>Balanced</i> , <i>Prefer backup</i> or <i>Prefer master</i> . By default, <i>Balanced</i> is selected.
Balanced	If you select this option, the leases are delivered to the clients by both servers equally.
Prefer backup	If you select this option, the leases are delivered to the clients by the backup server only.
Prefer master	If you select this option, the leases are delivered to the clients by the master server only.

- Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (one-to-one)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

DHCPv4 One-to-Many Smart Architecture

The [One-to-Many Smart Architecture](#), which is basically a star network topology of the DHCP servers of your choice, allows you to set up several failover channels between one master server and at least two backup servers to be used as backup. You actually can include as many servers as you want in this configuration as long as there are no power limitations or overload of the equipment managing the flow of information between the servers. This architecture also provides a shared management of the leases that you can configure according to your needs.

To configure a DHCPv4 One-to-Many smart architecture

- In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
- In the menu, select **+ Add > Server > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
- If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields according to the table below:

Table 24.5. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.

Managing DHCP Smart Architectures

Parameter	Description
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DHCP smart architecture**, select **One-to-Many**.

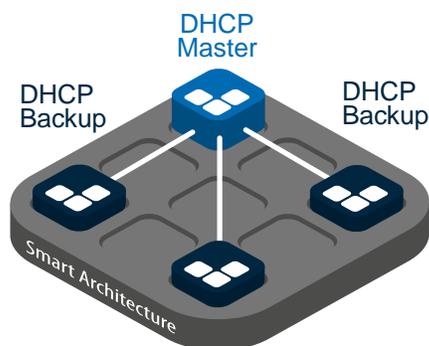


Figure 24.2. DHCPv4 One-to-Many smart architecture

7. Click on **NEXT**. The next page of the wizard opens.
8. In the list **Available DHCP servers**, you can select one by one the DHCP servers that you want to manage through the smart architecture. Ideally, you would configure at least three DHCP servers with this architecture.
9. Click on **+**. The selected server is moved to the list **Selected DHCP servers**. Repeat these actions as many times as needed.
10. Click on **NEXT**. The next page of the wizard opens.
11. In the drop-down list **Master DHCP server**, select the Master server in the smart architecture configuration.
12. Click on **NEXT**. The last page of the wizard opens.
13. In the list **DHCP peering assignment**, select the default failover channel named *Peering: failover-<smart_server_name> on DHCP ()*. Then, configure it following the table below:

Table 24.6. DHCPv4 One-to-Many failover parameters

Field	Description
Peering name	The default failover channel name, <i><failover-your.smart.server.name></i> , is displayed in the field. You can edit it.
Failover port	Type in the number of the port on the master server dedicated to the failover. By default, the port 847 is used, you can only use it once.
Failover peer port	Type in the number of the port on the backup servers dedicated to the failover. By default the port 647 is used, you can use it on each backup server if you want.
Automatic switch to partner-down delay (in hours)	Type in this field the amount of time (in hours) after which a failover channel in Communications-interrupted state should automatically switch to Partner-down. The accepted values are between 4 and 65535. By default, the option is disabled and the field is set to 0.
Peer DHCP server	Choose the DHCP backup server with which you want to configure the failover. By default, <i>None</i> is selected.
Split leases	In this drop-down list, you can choose how to split the leases between the two chosen servers: <i>Balanced</i> , <i>Prefer backup</i> or <i>Prefer master</i> . By default, <i>Balanced</i> is selected.
Balanced	If you select this option, the leases are delivered to the clients by both servers equally.
Prefer backup	If you select this option, the leases are delivered to the clients by the backup server only.
Prefer master	If you select this option, the leases are delivered to the clients by the master server only.

Click on **UPDATE** to commit your configuration. Your first failover channel is configured and present in the list **DHCP peering assignment** as such: *Peering: <failover_channel_name> on DHCP (<backup_server_name>)*.

Repeat this action in order to have a failover channel between the master and each backup server.

- Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (one-to-many)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

DHCPv4 Split-Scope Smart Architecture

The [Split-Scope Smart Architecture](#) allows you to distribute the management of ranges (and therefore leases) between two DHCP servers. They are set in an active/active configuration that ensures availability of the services at all times: if one server fails, the other can still lease IP addresses to the clients. You can actually choose the proportion of IP addresses (in percent) managed by each one of them.

To configure a DHCPv4 Split-Scope smart architecture

- In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
- In the menu, select **+ Add > Server > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
- If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields according to the table below:

Table 24.7. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The next page of the wizard opens.
- In the list **DHCP smart architecture**, select **Split-Scope**.

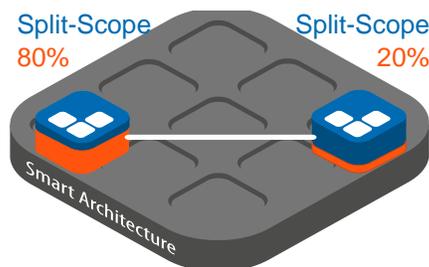


Figure 24.3. DHCPv4 Split-Scope smart architecture

- Click on **NEXT**. The next page of the wizard opens.
- In the list **Available DHCP servers**, you can select one by one the two DHCP servers if want to manage through the smart architecture.
- Click on **+**. The selected server is moved to the list **Selected DHCP servers**. Repeat these actions for the second server.
- Click on **NEXT**. The next page of the wizard opens.
- In the drop-down list **Master DHCP server**, select the Master server in the smart architecture configuration.

12. In the field **Distribution ratio (in percent)**, type in the ratio of IP ranges to be managed by the Master DHCP server you just selected. By default, 80 is proposed, meaning that the remaining 20% are listed and managed by the backup server.
13. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (split-scope)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

Note that a virtual failover channel is automatically created with the smart architecture, it is named *failover-<smart_architecture_name>* and listed on the page All failover channels.

DHCPv4 Single-Server Smart Architecture

The [Single-Server Smart Architecture](#) allows you to manage one single DHCP server that provides a backup. If the DHCP server crashes, the smart architecture configuration is saved and automatically applied to the new DHCP server managed through the Single-Server smart architecture.

To configure a DHCPv4 Single-Server smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 24.8. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page All servers.
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.

6. In the list **DHCP smart architecture**, select **Single-Server**.

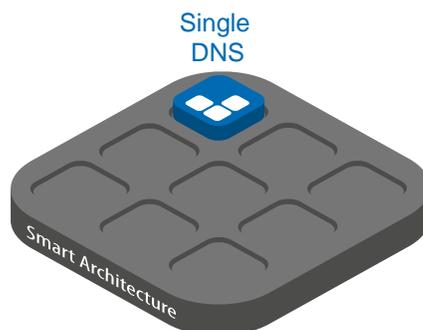


Figure 24.4. DHCPv4 Single-Server smart architecture

7. Click on **NEXT**. The next page of the wizard opens.
8. In the list **Available DHCP servers**, select the DHCP server that you want to manage through the smart architecture.
9. Click on **+**. The selected server is moved to the list **Selected DHCP servers**.
10. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (single-server)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button **☐** on the right-end side of the menu.

Note that a virtual failover channel is automatically created with the smart architecture, it is named *failover-<smart_architecture_name>* and listed on the page All failover channels.

Adding a DHCPv6 Smart Architecture

A smart architecture can manage IPv6 addresses and just like DHCPv4 can be configured without DHCP servers. Note that, with a DHCP v6 smart architecture, you still apply your configuration to a DHCP server managed on a SOLIDserver appliance running on an IPv4 address.

With DHCPv6, there are three different kinds of smart architectures: Single-Server, Split-Scope and Stateless. In the procedures below, we are going to describe the configuration of DHCPv6 smart architectures with DHCP servers but you can go through the configuration without adding any server and do it later. For more details, refer to the part [Adding a DHCP Server into a Smart Architecture](#).

DHCPv6 Single-Server Smart Architecture

The [Single-Server Smart Architecture](#) has the same advantages in DHCPv6 and DHCPv4, it allows you to manage one single DHCP server that provides a backup. If the DHCP server crashes, the smart architecture configuration is saved and automatically applied to the new DHCP server managed through the Single-Server smart architecture.

To configure a DHCPv6 Single-Server smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server (v6) > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.

- If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields according to the table below:

Table 24.9. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The next page of the wizard opens.
- In the list **DHCP smart architecture**, select **Single server**.

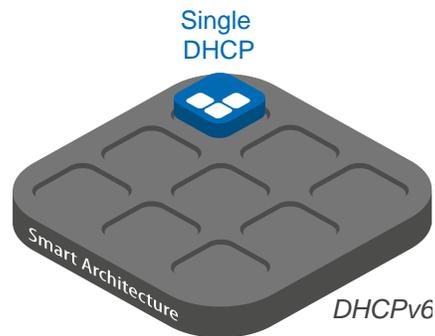


Figure 24.5. DHCPv6 Single-Server smart architecture

- Click on **NEXT**. The next page of the wizard opens.
- In the list **Available DHCP servers**, select the DHCPv6 server that you want to manage through the smart architecture.
- Click on **+**. The selected server is moved to the list **Selected DHCP servers**.
- Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (single-server)** in the column **Type**. You can

display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

Note that a virtual failover channel is automatically created with the smart architecture, it is named *failover-<smart_architecture_name>* and listed on the page All failover channels.

DHCPv6 Split-Scope Smart Architecture

The [Split-Scope Smart Architecture](#) allows you to distribute ranges of IP addresses between two DHCP servers. The active/active configuration ensures availability of the leasing service to clients. You can actually choose the proportion of IP addresses (in percent) managed by each one of them.

To configure a DHCPv6 Split-Scope smart architecture

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the menu, select **+ Add > Server (v6) > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on [NEXT](#). The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 24.10. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page All servers.
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on [NEXT](#). The next page of the wizard opens.
6. In the list **DHCP smart architecture**, select **Split-Scope**.

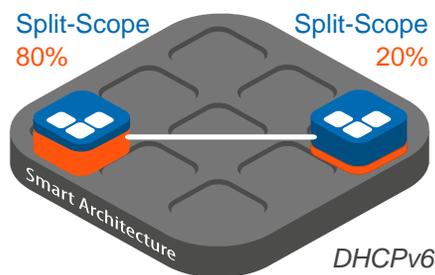


Figure 24.6. DHCPv6 Split-Scope smart architecture

7. Click on **NEXT**. The next page of the wizard opens.
8. In the list **Available DHCP servers**, you can select one by one the two DHCP servers if want to manage through the smart architecture.
9. Click on **+**. The selected server is moved to the list **Selected DHCP servers**. Repeat these actions for the second server.
10. Click on **NEXT**. The next page of the wizard opens.
11. In the drop-down list **Master DHCP server**, select the Master server in the smart architecture configuration.
12. In the field **Distribution ratio (in percent)**, type in the ratio of IP ranges to be managed by the Master DHCP server you just selected. By default, 80 is proposed, meaning that the remaining 20% are listed and managed by the backup server.
13. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (split-scope)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

Note that a virtual failover channel is automatically created with the smart architecture, it is named *failover-<smart_architecture_name>* and listed on the page All failover channels.

DHCPv6 Stateless Smart Architecture

The [Stateless Smart Architecture](#) allows you to set up a number of options to the scopes of the servers you choose to manage. The clients then have access to the options defined in the architecture. Keep in mind that there is no ranges, statics or leases management in a stateless architecture: you cannot create or provide them.

To configure a DHCPv6 Stateless smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server (v6) > DHCP smart architecture**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields according to the table below:

Table 24.11. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The next page of the wizard opens.
- In the list **DHCP smart architecture**, select **Stateless**.

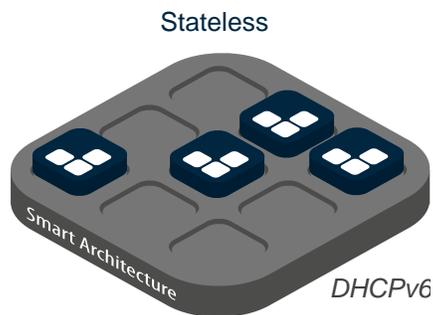


Figure 24.7. DHCPv6 Stateless smart architecture

- Click on **NEXT**. The next page of the wizard opens.
- In the list **Available DHCP servers**, you can select one by one as many DHCP servers as you want.
- Click on **+**. The selected server is moved to the list **Selected DHCP servers**. Repeat these actions as many times as needed.
- Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DHCP server and marked **Smart (stateless)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button **☐** on the right-end side of the menu.

Note that a virtual failover channel is automatically created with the smart architecture, it is named *failover-`<smart_architecture_name>`* and listed on the page *All failover channels*.

Editing a DHCP Smart Architecture

Once created, you can edit a smart architecture to change the servers it manages, edit the server roles, change the smart architecture type or convert a server into a smart architecture.

Adding a DHCP Server into a Smart Architecture

Once a smart architecture is properly configured and applied, you can add DHCP servers whenever you want. First, to add a DHCP server, refer to the section [Managing DHCP Servers](#). According to the DHCP smart architecture chosen, if you do not complete the architecture with all the necessary servers, the smart architecture may not work properly. Make sure that you have added all the necessary DHCP servers into the smart architecture.

When you add one or more DHCP servers into a smart architecture, the smart data is automatically replicated from the architecture to the DHCP servers it manages. So if the smart architecture is empty (first use), the DHCP server added is totally overwritten.

To add a DHCP server into DHCP smart architecture

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on [EDIT](#). The wizard **Edit a DHCP server** opens.
4. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on [NEXT](#). The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. In the list **DHCP server type**, make sure *DHCP smart architecture* is selected. Click on [NEXT](#). The **Manage a DHCP server** page opens.
6. If need be, modify the smart architecture basic parameters according to the table below:

Table 24.12. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .

Parameter	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **[NEXT]**. The next page of the wizard opens.
- In the list **DHCP smart architecture**, modify the type of your DHCP smart architecture if need be. Click on **[NEXT]**. The next page of the wizard opens.
- In the list **Available DHCP servers**, select a server to add in the smart architecture and click on **[+]**. The server has been moved to the list **Selected DHCP servers**. Repeat this action for as many servers as needed. You can remove any of them from the selected servers list by clicking on **[-]**.
- For a Single-Server smart architecture, go to the last step of this procedure. Otherwise, click on **[NEXT]**. The next page of the wizard opens.
- In the drop-down list **Master DHCP server**, edit the master server if need be.
- For a Split-Scope architecture, in the field **Distribution ratio (in percent)**, type in the ratio of IP ranges to be managed by the selected Master DHCP server, the rate is managed by the backup server.
- If need be, edit the existing failover ports and split leases parameters between the master and backup servers.
- Click on **[OK]** to complete the operation. The report opens and closes. You can display or hide the physical servers managed through your smart architecture using the button **[☰]** on the right-end side of the menu. The **DHCP Smart members** column of the smart architecture displays the name of the new master server between brackets next to the name of the other backup servers.

Removing a DHCP Server from a Smart Architecture

Whenever you want, you can remove one or more DHCP servers from a DHCP smart architecture. When you remove one, the configuration applied on this server is conserved on the DHCP server previously removed.

To remove a DHCP server from a smart architecture

- In the sidebar, go to **🔍 DHCP > Servers**. The page **All servers** opens.
- At the end of the line of the smart architecture of your choice, click on **[⚙️]**. The properties page opens.
- In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DHCP server** opens.
- If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Click on **[NEXT]**. The page **Manage a DHCP server** opens.
- Click on **[NEXT]**. The next page of the wizard opens.
- Click on **[NEXT]**. The next page of the wizard opens.

8. The servers managed by the smart architecture are listed in the list **Selected DHCP servers**. You can remove any of them by clicking on . The server(s) is moved to the list **Available DHCP servers**.
9. For a Single-Server smart architecture, go to the last step of this procedure. Otherwise, click on **NEXT**. The next page of the wizard opens.
10. If the smart architecture is still managing servers: in the list **Master DHCP server**, change the master server if need be. Click on **NEXT**. The next page of the wizard opens.
11. If the smart architecture is still managing servers: modify the failover ports on each server and/or the split leases parameters if need be.
12. Click on **OK** to complete the operation. The report opens and closes. The servers that have been removed are listed as DHCP servers of whatever kind in the list **Type**. If your smart architecture is still managing physical servers, you can display or hide them using the button  on the right-end side of the menu.

Changing the DHCP Server Roles within a Smart Architecture

You can change the role of DHCP servers within a smart architecture. For instance, you can change a master server into a slave server within a One-to-One smart architecture at any given time.

To change the role of DHCP servers within a smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DHCP server** opens.
4. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Click on **NEXT**. The page **Manage a DHCP server** opens.
6. Click on **NEXT**. The next page of the wizard opens.
7. Click on **NEXT**. The next page of the wizard opens.
8. The servers managed by the smart architecture are listed in the list **Selected DHCP servers**. You can remove any of them and add a new one by clicking on  or . The server(s) is moved accordingly between the lists **Selected DHCP servers** and **Available DHCP servers**.
9. For a Single-Server smart architecture, go to the last step of this procedure. Otherwise, click on **NEXT**. The next page of the wizard opens.
10. In the drop-down list **Master DHCP server**, select the master server.
11. For a Split-Scope architecture, in the field **Distribution ratio (in percent)**, type in the ratio of IP ranges to be managed by the selected Master DHCP server, the rate is managed by the backup server.
12. If need be, edit the existing failover ports and split leases parameters between the master and backup servers.

- Click on **OK** to complete the operation. The report opens and closes. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu. The column **DHCP Smart members** of the smart architecture displays the name of the new master server between brackets next to the name of the other backup servers.

Changing the Type of a DHCP Smart Architecture

The type of a DHCP smart architecture can be easily changed while keeping all DHCP configuration and data you already set. For instance, you already have a DHCP smart architecture configured in One-to-One that includes two DHCP servers -one in master and the other in slave- and you plan to change your smart architecture type into Split-Scope. By editing the smart architecture, you can change its type and configure the role of servers.

To edit the type of a DHCP smart architecture

- In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
- At the end of the line of the smart architecture of your choice, click on . The properties page opens.
- In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DHCP server** opens.
- If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the list **DHCP server type**, make sure *DHCP smart architecture* is selected. Click on **NEXT**. The **Manage a DHCP server** page opens.
- If need be, modify the smart architecture basic parameters according to the table below:

Table 24.13. DHCP smart architecture basic parameters

Parameter	Description
DHCP server name	Name your server with a valid FQDN.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DNS). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any configuration set via the drop-down list <i>Advanced Properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The next page of the wizard opens.

8. In the list **DHCP smart architecture**, select a different smart architecture type.
9. Click on **[NEXT]**. The next page of the wizard opens.
10. If your smart architecture manages servers, they are listed in the list **Selected DHCP servers**. You can select them one by one and click on **[+]**. They are moved to the list **Selected DHCP servers**.
11. For a *Single-Server* smart architecture, go to the last step of this procedure.
12. For any other smart architecture type, click on **[NEXT]**. The next page of the wizard opens.
 - a. In the drop-down list **Master DHCP server**, select the Master server in the smart architecture configuration.
 - b. For a Split-Scope architecture, in the field **Distribution ratio (in percent)**, type in the ratio of IP ranges to be managed by the selected Master DHCP server. The rest is managed by the backup server.
 - c. If need be, edit the existing failover ports and split leases parameters between the master and backup servers. For more details, refer to the relevant procedure of the section [Adding a DHCPv4 Smart Architecture](#) or [Adding a DHCPv6 Smart Architecture](#).
13. Click on **[OK]** to complete the operation. The report opens and closes. The page **All servers** is visible again. The column **Type** displays the modification you performed on the smart architecture.

Converting a DHCP Server into a Smart Architecture

To keep a server configuration and avoid configuring a smart architecture to match the server settings before adding it into the smart, you can convert DHCP servers into smart architectures.

Keep in mind that once you converted a DHCP server into a smart, it is no longer listed on the page *All servers*. You have to add it again to be able to manage it, on its own or from a smart architecture.

During the conversion, you can add DHCP servers into the smart architecture. Considering that you might want to manage the server you converted from the smart architecture, we recommend converting the server and then editing the smart to add the servers as detailed in the section [Adding a DHCP Server into a Smart Architecture](#).

To convert a DHCP server into a smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Right-click over the **Name** of the server of your choice. The contextual menu appears.
3. Click on **[E] Edit**. The wizard **Edit a DHCP server** opens.
4. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. In the list **DHCP server type**, select *DHCP smart architecture*.
6. Click on **[NEXT]**. The next page of the wizard opens, it displays the server details.
7. Click on **[NEXT]**. The next page of the wizard opens.

8. In the field **DHCP smart architecture**, select the DHCP or DHCPv6 smart architecture of your choice.

For a conversion to Single-Server smart architecture, go to the last step of this procedure.

9. Click on **NEXT** until the last page of the wizard opens.
 - a. For a conversion to One-to-One smart architecture, you can configure the failover channel with a specific *Failover port*, *Failover peer port*, *Automatic switch to partner-down delay* and *Split lease* distribution.
 - b. For a conversion to One-to-Many smart architecture, you can configure the failover channels listed in the field *DHCP peering assignments*. Select them one by one to load their details in the fields and click on **UPDATE** to save your changes.

If you converted a Microsoft server, you cannot have more than 31 failover channels listed. Any extra channel is ignored.

- c. For a conversion to Split-Scope smart architecture, you can select a *Master DHCP server* and set the *Distribution ratio (in percent)* of the Master.

For more details regarding the smart architecture configuration, refer to the section [Adding a DHCPv4 Smart Architecture](#).

10. Click on **OK** to complete the operation. The report opens and closes. In the column **Type**, the server is now listed as a smart architecture.

Handling the Status Locked Synchronization

SOLIDserver provides a consistency check for the smart architectures. Once you configured a smart architecture with the server(s) you want to manage, the smart configuration is checked before it is sent to the physical server(s): this ensures the consistency of the configuration and avoids pushing useless information to the server:

- If the check is conclusive, the information is sent to the server and, on the page *All servers*, its status is *Synchronized*.
- If any error is found, the verification stops and the server **Sync** status changes to **Locked Synchronization** once the page is refreshed. To get a valid synchronization status again, you need to "undo" the latest changes. This action loads a new synchronization and uploads the status accordingly.

Once the server is in *Locked synchronization*, the corrupted configuration file is automatically stored locally on the appliance and available for download in the Local files listing. It is named `<server_name>-dhcpd.conf`. We advise that you take a look at this file because after the first found error, the check stops and returns the *Locked synchronization* status. So if there are several errors, the status is returned over and over again until the file is conclusive and can be sent to the physical server.

The check for failure in the configuration file can be done through CLI (we recommend it) or from the GUI.

To check for failure in a DHCPv4 configuration file through CLI

1. Open an SSH session.
2. Use the following command to retrieve the list of corrupted files:

```
# ls -la /data1/exports/*-dhcpd.conf
```

3. Use the following command to get a precise list of all the errors:

```
# /usr/local/nessy2/bin/dhcpd -t -4 -cf /data1/exports/<server_name>-dhcpd.conf
```

4. Adjust identified statements, once the check runs again, the *Locked Synchronization* status disappears if you now have a valid configuration.

To check for failure in a DHCPv6 configuration file through CLI

1. Open an SSH session.
2. Use the following command to retrieve the list of corrupted files:

```
# ls -la /data1/exports/*-dhcpd6.conf
```

3. Use the following command to get a precise list of all the errors:

```
# /usr/local/nessy2/bin/dhcpd -t -6 -q -cf /data1/exports/<server_name>-dhcpd6.conf
```

4. Adjust identified statements, once the check runs again, the *Locked Synchronization* status disappears if you now have a valid configuration.

To look for DHCP errors on the page Syslog of the local appliance

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
3. In the SOLIDserver drop-down list, verify that the local appliance is selected. Only the host-name appears with no IP address.
4. In the Services filed, select *dhcpd*. The logs appear.

Deleting a DHCP Smart Architecture

At any time, you can decide to stop managing your DHCPv4 or DHCPv6 servers through the smart architectures. You might not need to delete a smart architecture, editing it might be enough. For more details, refer to the section [Changing the Type of a DHCP Smart Architecture](#).

Before deleting a smart architecture, keep in mind that:

- Deleting a smart architecture does not delete any data from the physical server: it means that you stop managing the server via the smart architecture. However, **the configuration backup of the smart architecture is deleted**. Therefore, if the server crashes after the smart architecture deletion, you have to configure everything again manually.
- **You cannot delete a smart architecture if it is still managing DHCP servers.**

To delete a DHCP smart architecture

1. In the sidebar, go to  **DHCP > Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. If the smart architecture is managing DHCP servers, remove them according to the [Removing a DHCP Server from a Smart Architecture](#) section.
4. Tick the smart architecture(s) you want to delete.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.

6. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is no longer listed on the page **All servers**.

Defining a DHCP Smart Architecture as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a smart architecture as one of the resources of a specific group allows the users of that group to manage the architecture in question as long as they have the corresponding rights granted.

Granting access to a smart architecture as a resource also grants access to every physical server it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 25. Managing DHCP Servers

Within the DHCP module, the server is the highest level of the hierarchy where you set the basis of any DHCP configuration. You can either manage servers independently or within a smart architecture that provides a backup of your configuration and a dedicated failover between a master server and its backup(s) servers. For more details regarding the available smart architectures for DHCPv4 or DHCPv6, refer to the chapters [Deploying DHCP Smart Architectures](#) and [Managing DHCP Smart Architectures](#).

In IPv4, you can create EfficientIP DHCP, EfficientIP DHCP Package, Microsoft DHCP and Nominum DCS DHCP servers on the page *All servers*. In IPv6, you can only create EfficientIP DHCP servers.

Whether you manage IPv4 or IPv6 addressing, the IP address of the DHCP servers you manage must be in IPv4 and belong to one of the networks that your configured part of the networks you configured, to make sure it can provide clients with IP addresses. If you manage large networks, DHCP servers can rely on DHCP relays (also called helpers), for more details refer to the section [DHCP Relay Agents](#).

Browsing DHCP Servers

To put it simply, the server is a container for all the information necessary to provide IP addresses to the DHCP clients. Keep in mind that any parameter and/or option set at a lower level overwrites any configuration set at server level.

These servers can be configured to provide IPv4 and IPv6 addresses, obviously the options available change from one version to the other as in essence, DHCPv4 and DHCPv6 protocols can be considered to be two different protocols although they serve the common goal of providing the addresses to DHCP clients. Both versions of the Dynamic Host configuration Protocol allow to configure the server and provide either dynamic addressing or fixed addressing: in the image below the two branches of the tree symbolize both types of addressing. On the left are represented the level of hierarchy necessary to set up dynamic addressing and on the right the fixed addressing.

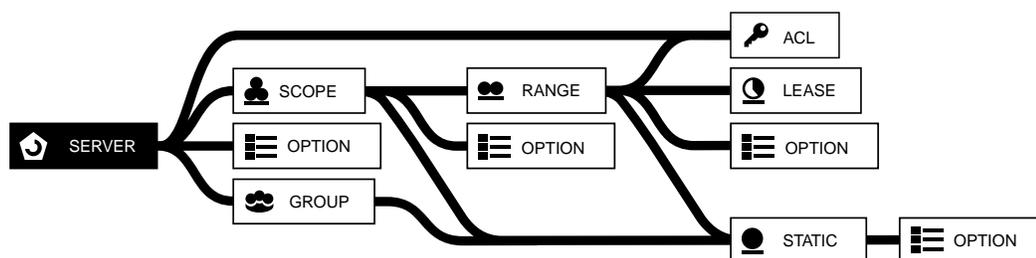


Figure 25.1. The server in the DHCP hierarchy

So, in short, the very first step of the DHCP implementation is the creation of a server with a unique IP address within which you must create at least one scope that listens on a particular part of the network and discover clients' request and answer them at the best of its capacity. Afterward, you decide to set up dynamic and/or fixed addressing for the DHCP clients.

Browsing the DHCP Servers Database

To display the list of DHCP servers

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. To list the IPv4 DHCP servers:
 - a. In the column **Protocol**, right-click over *IPv4*. The contextual menu.
 - b. Click on . Only the IPv4 DHCP servers are listed.
3. To list the IPv6 DHCP servers:
 - a. In the column **Protocol**, right-click over *IPv6*. The contextual menu.
 - b. Click on . Only the IPv6 DHCP servers are listed.

To display a DHCP server properties page

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. Filter the list if need be.
3. At the end of the line of the server of your choice, click on . The properties pages opens.

Customizing the Display on the Page All Servers

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the DHCP Server Statuses

The column **Status** provides information regarding the servers you manage.

Table 25.1. DHCP server statuses

Status	Description
 <i>OK</i>	The server is configured
 <i>Unknown</i>	The server does not have a status as it has not synchronized yet.
 <i>Timeout</i>	The server does not answer anymore due to a scheduled configuration of the server.
 <i>Invalid credentials</i>	The SSL credentials are invalid or the server is already managed by another appliance and you need to specify your credentials again. For more details, refer to the section Editing a DHCP Server .
 <i>Syntax error</i>	The server configuration could not be parsed properly.
 <i>License</i>	The license used in SOLIDserver is not compliant with the added server: the license is invalid.
 <i>Invalid settings</i>	There was a setting error during the server declaration. For instance, some settings were added to a server that does not support them or a smart architecture is not managing any physical server.
 <i>Insufficient privileges</i>	The account used to add the Agentless DHCP server does not have sufficient privileges to manage it.

Status	Description
🔴 <i>Locked synchronization</i>	The server configuration is not viable. For more details, refer to the section Handling the Status Locked Synchronization .

Note that the column **Sync** column changes in accordance with the column **Status**. While the server synchronization is not 🟢 OK yet, the column *Sync* might be 🟡 Busy; it can also be in *Locked synchronization*.

The column **Multi-status** provides you with emergency, warning, critical, error or informational messages regarding the server compatibility with Hybrid. For more details, refer to the section [Understanding the Column Multi-Status](#).

Managing EfficientIP DHCP Servers

The DHCP management module supports several EfficientIP DHCP servers, including the DHCP server embedded in SOLIDserver device and EfficientIP DHCP package running the ISC DHCP server on Linux.

Configuring the Listening Network Interfaces

The DHCP server selects the listening network interfaces via the DHCP scopes¹. To make the server listen on an interface, you have to create a scope that includes one or several local interfaces of the DHCP server to allow the server to reply to the DHCP client requests.

For instance, your DHCP server has 2 network interfaces configured: *192.168.10.3* and *192.168.10.5*. To listen to both interfaces, you have to configure a scope with the network address *192.168.10.0* and the netmask *255.255.255.0*. For more details regarding scope management, refer to the chapter [Managing DHCP Scopes](#).

EfficientIP DHCP server implements the safe DHCP failover protocol. For more details, refer to the chapter [Managing Failover Channels](#).

Adding an EfficientIP DHCP Server

From the page *All servers*, you can add an EfficientIP DNS Server to manage its configuration, all its data and monitor it. Before adding the server, keep in mind that:

- EfficientIP DHCP servers can provide IPv4 or IPv6 addressing.
- **The SNMP protocol is no longer supported as managing protocol for a server.** Therefore:
 - You can no longer add a server managed via SNMP.
 - EfficientIP DHCP servers prior to version 4.0.x are no longer supported.
 - Your existing servers in version 4.0.x or prior, migrated to 7.0, are still managed via SNMP and listed in the GUI. However, the management of these servers is not detailed in this guide. For more details, refer to the guide *SOLIDserver-Administrator-Guide-5.0.4.pdf*.
- **SSL is used to manage a server while SNMP is used to monitor it.** Therefore:
 - If you manage a DHCP server in version 4.0.x legacy, editing it automatically changes the management protocol to use SSL instead of SNMP. This operation is **non-reversible**.
 - You can configure the SNMP monitoring parameters of the server. For more details, refer to the section [Editing the SNMP Monitoring Parameters of an EfficientIP DHCP Server](#).

¹If you do not set a listening scope, you must configure a relay to communicate with DHCP clients.

- **EfficientIP DHCP servers manage the IPv4 static reservations like leases.** The MAC address specified during the static reservation identifies the clients' IP address and allocates it a lease as well as soon as they are visible on the network. Once the lease is allocated, if the IPAM and DNS replication are configured, the data is sent to the IPAM and creates the corresponding DNS entries. For more details, refer to the section [Adding DHCPv4 Statics](#).
- **You should not set a VIP as management address of your DHCP server.**
- Since version 7.0, a random password is generated when you add a server. The default SSH credentials of the account *admin* are no longer used to manage the server but to generate this random password.

For servers added before the upgrade to version 7.0, switching to this new management system is not automatic. You need to edit the servers. For more details, refer to the section [Editing a DHCP Server](#).

To add an EfficientIP DHCP server

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the menu, select **+ Add > Server** or **Server (v6) > EfficientIP DHCP**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the following fields to set up the basic server configuration:

Table 25.2. DHCP server basic parameters

Field	Description
DHCP server name	In this field, fill in a FQDN name for your server. This field is required.
Management IP address	In this field, fill in the IPv4 address of your server. This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DNS. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> . This field is optional.

5. If you have modified the SSH password or if the server is already managed by another appliance, tick the box **Configure enrollment parameters**. If not, go to step 7.

Once you have ticked the box, the field **"Admin" account password** appears. The default *Admin* account password is automatically filled.

6. In the field **"Admin" account password**, enter your SSH password.
7. If you want to edit the server SNMP parameters², tick the box **Configure SNMP monitoring parameters**. If not, go to step 8.

²The SNMP protocol parameters are used to monitor and retrieve the server statistics.

Once you ticked the box, the following fields appear:

Table 25.3. SNMP parameters used to monitor the server statistics

Field	Description
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
Use TCP	Tick the box if you want to use the TCP protocol instead of the UDP when the network link is not reliable.
SNMP profile	The SNMP profile used to retrieve the statistics. By default, <i>standard v2c</i> is selected. The list contains the default profiles (<i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i>) and the ones you may have created. Each profile has its own level of security and enables the definition of a global security policy. For more details, refer to the section Managing SNMP Profiles .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.

- Depending on the administrator configuration, you may be able to configure advanced properties according to the table below.

Table 25.4. Advanced properties configuration

Field	Description
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The server is listed.

To edit the SNMP monitoring parameters on an existing server, refer to the section [Editing the SNMP Monitoring Parameters of an EfficientIP DHCP Server](#).

Editing the SNMP Monitoring Parameters of an EfficientIP DHCP Server

Once added to the page *All servers*, you can edit the SNMP monitoring parameters of an EfficientIP DHCP server.

To edit the SNMP monitoring parameters of an EfficientIP DHCP server

- In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
- At the end of the line of the server or smart architecture of your choice, click on . The properties pages opens.
- In the panel **SNMP monitoring parameters**, click on . The wizard **SNMP parameters** opens.

4. Edit the SNMP parameters according to your needs:

Table 25.5. SNMP parameters used to monitor the server statistics

Field	Description
SNMP version	The version of the SNMP protocol used to retrieve the statistics. It can be either v1, v2c or v3. By default, v2c is selected.
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.
Use bulk	If you use SNMP version 2 or 3, you can choose to use a bulk transfer of data. This compact SNMP request method accelerates transfers by sending several requests at once. By default, it is set to Yes.
Use TCP	Choose to use the TCP protocol instead of the UDP when the network link is not reliable. By default UDP is used, the drop-down list is set to No.

5. Click on **[NEXT]**. The page **SNMP profile** opens.
6. In the drop-down list **SNMP profile**, choose a profile using the same version of the SNMP protocol as the one you selected in the field *SNMP version*.

If you created SNMP profiles, you can choose one of your profiles. They are listed only if they use the same version of the SNMP protocol as the one you selected on the previous page.

Note that the SNMP profiles you can choose from must be configured on the appliance you are currently working with. For more details, refer to the section [Managing SNMP Profiles](#).

7. Click on **[OK]** to complete the operation. The report opens and closes. The changes are listed in the panel.

Managing Agentless Microsoft DHCP Servers

You can add Microsoft DHCP servers to manage them from the page *All servers*. They can be included into an Active Directory (AD) domain or not.

Once you manage a Microsoft server, you can also manage its scopes, ranges, statics, leases. Depending on the version of Windows and if you add it to a smart architecture, you can manage its failover relationships.

The management of Microsoft DHCP servers is based on Microsoft Remote Procedure Calls (MSRPC) and allows to retrieve and display data in real-time and avoid installing any WinDHCP agent. Microsoft DHCP servers with agent are not supported.

Understanding the Range Management on Microsoft Servers

The way to manage the ranges within SOLIDserver and Microsoft DHCP server is different. You can create as many ranges as you need with SOLIDserver but only one with Microsoft DHCP. So, when SOLIDserver overwrites Microsoft DHCP server configuration:

- The unique Microsoft range start and end addresses match the start and end addresses of a SOLIDserver scope.

- A number of exclusion ranges are created on the Microsoft server to match the ranges you created with SOLIDserver.

This mechanism allows SOLIDserver to offer the same services as Microsoft but it displays them differently than Windows Administrative Tools of your Microsoft DHCP server. The ranges that you create with SOLIDserver correspond to a unique range with a number of exclusion ranges.

SOLIDserver DHCP configuration vs. Microsoft DHCP configuration

With SOLIDserver when you create a scope with the start address 192.168.10.0 and the end address 192.168.10.255, the configuration is pushed onto Microsoft DHCP server exactly the same. However, the way to deal with the ranges differ.

When you create the two following ranges with **SOLIDserver**:

First range

192.168.10.5 - 192.168.10.10

Second range

192.168.10.25 - 192.168.10.100

The configuration looks as follows in the **Microsoft DHCP configuration**:

One unique range

192.168.10.0 - 192.168.10.255. It basically corresponds to the scope start and end addresses.

Three exclusion ranges

192.168.10.1 - 192.168.10.4

192.168.10.11 - 192.168.10.24

192.168.10.101 - 192.168.10.254

Prerequisites

- A Microsoft Windows Server 2008, 2008 R2, 2012 R2, 2016 or 2019. The server must:
 - Have the TCP ports 135 and 445 open. They are used by the port mapper interface, the service that indicates to the clients which port provides access to each service.
 - Have Firewall policies that allow traffic between SOLIDserver and the Microsoft servers it manages.
 - In Windows Server 2008, RPC uses by default the dynamic port range 49152-65535. Note that you can reduce the number of available ports, using netsh, as long as the range of ports contains at least 255 ports³.
- The credentials of a member of the group *DHCP Administrators*. Users with insufficient privileges cannot manage the server.
- The service DHCP must be properly started. For more details, refer to the chapter [Configuring the Services](#).

Limitations

The management of Microsoft DHCP servers within SOLIDserver has some limitations. For more details regarding the Microsoft limitations, refer to their documentation.

³For information, refer to <http://support.microsoft.com/kb/929851> .

Server Limitations

- To add or manage Microsoft servers from SOLIDserver, a user must have administrator rights over the Microsoft server in your Windows environment.
- To display Microsoft servers in SOLIDserver, a user must have reading rights over the Microsoft server in your Windows environment.
- You cannot manage Microsoft servers that are Master in one configuration and backup in another on your Windows environment. Once added to the GUI, they take on one role or the other.
- Changes performed directly on the Microsoft server are not automatically transferred to SOLIDserver. You must select the server on the page *All servers* and synchronize it via the menu *Edit*.
- The synchronization of a Microsoft server within SOLIDserver does not work if the server is managed via a smart architecture, the smart configuration overwrites the new data.
- Microsoft policies are not supported, any policy configured on a Microsoft server is ignored.
- The statistics of Microsoft servers are not retrieved, the page *Analytics* does not include them.

DHCP Options Limitations

- Encapsulated DHCP options are not supported by Microsoft DHCP servers.

Lease Limitations

- The start date of a lease is unknown. SOLIDserver displays an arbitrary start date that corresponds to the moment when the lease is detected.
- DHCP configurations involving a very large number of leases trigger refresh problems. By default, the registry database entry *module.dhcp.refresh_server_time* refreshes leases every 10 seconds, when there are a lot of leases it can overload the service and create a loop. To avoid this problem, you need to increase the value of the registry entry.

ACL Limitations

- You cannot configure ACL on Microsoft servers.

Adding an Agentless Microsoft DHCP Server

Once you reviewed the [Prerequisites](#) and [Limitations](#), you can add an Agentless Microsoft DHCP server to manage its scopes, ranges, statics and leases.

If your Microsoft DHCP server is integrated to an AD with several forests, you can use the *Expert mode* during the server addition to specify the AD domain you want to authenticate.

If you manage an Agentless Microsoft DHCP server 2012 R2 or higher from a smart architecture, you can manage its failover relationships. For more details, refer in the section [Managing the Failover Channels of an Agentless Microsoft DHCP Server](#).

To add an agentless Microsoft DHCP server

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > Microsoft DHCP**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the following fields to set up the basic server configuration:

Table 25.6. Microsoft DHCP server basic parameters

Field	Description
DHCP server name	Type in an FQDN name for the server. This field is required.
Management IP address	Type in the IPv4 address of the Microsoft DHCP server you want to manage ^a . This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DNS. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> . This field is optional.

^aWith the proper [Configuring the Network](#), if you enter the name of your DHCP server in this field and click on **SEARCH**, the IP address is retrieved from the DNS and displayed.

5. Click on **NEXT**. The last page of the wizard opens.
6. In the field **Login**, type in the name of user with sufficient managing privileges over the Windows DHCP server.
7. In the field **Password**, type in the corresponding password.
8. If your Microsoft DHCP server is integrated to Active Directory and contains several forests:
 - a. Tick the box **Expert mode (AD)**. The field *Domain* appears
 - b. In the field **Domain**, type in the full domain or subdomain.
9. Depending on the administrator configuration, you may be able to configure advanced properties according to the table below.

Table 25.7. Advanced properties configuration

Field	Description
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

10. Click on **OK** to complete the operation. The report opens and closes. The server is listed, the column **Version** indicates the Microsoft server version.

Managing the Failover Channels of an Agentless Microsoft DHCP Server

Within SOLIDserver, Microsoft failover relationships are called failover channels.

To manage Microsoft failover channels, you must manage Microsoft servers from a smart architecture. As managing servers from a smart erases their content, you must either prepare the smart architecture before managing server with it or convert the server into a smart architecture to keep its configuration.

Once you manage a Microsoft server with a smart architecture, you can even associate its scopes with a specific failover. For more details, refer to the section [Defining a Specific Failover Channel for a Scope](#).

Prerequisites

- Meeting Microsoft [Prerequisites](#).
- Microsoft DHCP servers 2012 R2 or higher.
- All the Microsoft servers you manage in a smart architecture must use the same NTP pool.

Limitations

- Taking into account Microsoft [Limitations](#).
- Within SOLIDserver a failover is between two servers.
- SOLIDserver cannot create failover channels for Microsoft servers if they do not manage valid ranges. Without a valid range, no failover can be created from SOLIDserver and pushed to the server.
- Only the failover mode *Load balance* is supported:
 - Three ratios are supported *100-0* (Prefer Master), *0-100* (Prefer Backup) and *50-50* (Balanced). Any other ratio is overwritten by the closest ratio we support.
 - If you manage failovers in mode *Hot standby*, their mode is switched to *Load balance*.
- If you manage Microsoft servers from a smart architecture:
 - It must only manage Microsoft servers.
 - It must be One-to-One or One-to-Many smart architecture. You cannot manage Microsoft failovers from a Single-Server smart architecture.
 - In One-to-Many smart architectures, you can only have one *Master* server in the smart architecture. It takes on the role of *primary* server.
 - In One-to-Many smart architectures, you cannot manage more than 31 failover channels, i.e. 32 Microsoft servers.
 - The unique port numbers that you must specify in the GUI are actually ignored. Microsoft automatically sets the proper port for the communication.
- If you manage Microsoft servers outside a smart architecture:
 - You can manage their content but not their failover channels. You can only display them on the page *All failover channels*.
 - You cannot configure advanced properties between the IPAM and the DHCP as they rely on failover channels. For more details, refer to the chapter [Managing Advanced Properties](#).

To manage your Microsoft failover channels from a smart architecture we recommend to:

1. **Add all the Microsoft DHCP servers that you want to manage from SOLIDserver.**

For more details, refer to the section [Adding an Agentless Microsoft DHCP Server](#).

2. **Convert one Microsoft server into a One-to-One or One-to-Many smart architecture.** We recommend converting the Master server as it contains all the configuration data. During the conversion, you can configure the failover channels:

- On One-to-One smart architectures, you cannot have more than one failover channel.
- On One-to-Many smart architectures, you cannot have more than 31 failover channels.

For more details, refer to the section [Converting a DHCP Server into a Smart Architecture](#).

3. **Add again the server you converted into a smart**, if you want to manage it and not only use its configuration.

For more details, refer to the section [Adding an Agentless Microsoft DHCP Server](#).

4. **Edit the smart architecture to make it manage all the Microsoft servers** and specify the Master.

For more details, refer to the section [Adding a DHCP Server into a Smart Architecture](#).

Managing ISC DHCP Servers

Efficient IP provides its software versions through native packages of operating system. Installing a DHCP package allows you to use the DHCP module of SOLIDserver at the best of its potential on Linux: it allows you to manage an ISC server through an EfficientIP DHCP server and benefit from all the options that come with it (statistics retrieved via SNMP...).

The addition of Linux Packages v4 that respected the SNMP protocol and worked with SSL is not supported. If you migrated your database with such servers, refer to the guide *SOLIDserver-Administrator-Guide-5.0.4.pdf* available on our website ⁴.

Managing EfficientIP ISC Linux Packages

In the sections below are a set of procedures to successfully install the DHCP packages, formerly known as EfficientIP ISC Linux Packages v5, on Linux Debian and CentOS/RedHat.

Installing the EfficientIP DHCP Package for Linux Debian 9 and prior - 64 bits

You must take into account the [prerequisites](#) before [installing](#) a DHCP Debian package.

Prerequisites

- The DHCP package file, *ipmdhcpxx-y.y.y-debianxx-amd64.deb*, whose name provides you with a number of information separated by hyphens: the type of package (*ipmdhcpxx*: a DHCP package with a DHCP in version *xx* where *xx* is *x* dot *x*), the version of SOLIDserver (*y.y.y*); the version of Debian (*debianxx* where *xx* is *x* dot *x*) and finally the Debian architecture (*amd64*).

In the procedure below, this file is referred to as `<ipmdhcpxx-y.y.y-debianxx-amd64.deb>`.

- The EfficientIP ISC package platform must have at least 20 Mb of free disk space.

⁴At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation* and in the folder */docs/5.0.4*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

- The EfficientIP ISC package may need certain libraries of your operating system, you must have a *shell* access with *root* login in local, via *ssh*⁵ on the server to be installed.
- You must make sure that no other DNS/DHCP service on your Linux is running : it would interfere with the BIND/ISC package installation.
- You must make sure that SOLIDserver and Debian are set to the same time and date.
- You must make sure that Apache server is up-to-date.
- You must make sure that the service *dbus* is installed.
- You must make sure that HTTPS (port 443), the DHCP service (port 67) and the failover ports (647-667 and 847-867) are not blocked by a network filtering process (firewall).

If your Apache configuration already uses the port 443, you have to create an additional IP-based VirtualHost dedicated to the DNS/DHCP management.

Installing the EfficientIP DHCP Package

You can install the EfficientIP DHCP Package on Debian Linux.

If you have not installed the DNS packages yet, you need to:

1. Follow the procedure [To install the EfficientIP DHCP Package on Debian](#).
2. Follow the procedure [To complete the DHCP package installation on Debian if the DNS package is not installed](#).

If you already installed the DNS packages, you only need to follow the procedure [To install the EfficientIP DHCP Package on Debian](#) below.

The installation procedure below includes the commands that make the web services configurable.

To install the EfficientIP DHCP Package on Debian

1. Open an SSH session.
2. Stop and disable your DHCP software, using the following commands:

```
# service isc-dhcp-server stop
# update-rc.d -f isc-dhcp-server remove
```

3. Install the dependency packages, ONLY if you have not installed the EfficientIP DNS package, using the following commands:

- a. For Debian 8 and prior:

```
# apt-get install php5
# apt-get install sudo
# apt-get install snmpd
# apt-get install sqlite
# apt-get install php5-sqlite
```

- b. For Debian 9:

```
# apt-get install php
# apt-get install sudo
# apt-get install snmpd
# apt-get install sqlite
# apt-get install php-sqlite3
```

⁵You could also connect via *telnet* but, for security purposes, we recommend that you favor *ssh*.

4. Install the EfficientIP DHCP package, using the following command:

```
# dpkg -i <ipmdhcpxx-y.y.y-debianxx-amd64.deb>
```

5. Make the web services configurable: in the directory `/etc/sudoers.d`, create the file `ipmdhcp` containing the line below.

```
www-data ALL = NOPASSWD: /usr/local/nessy2/script/install_dhcpd_conf.sh, \  
/usr/local/nessy2/script/install_dhcpd6_conf.sh
```

6. Set the users access rights as follows:

```
# chmod 440 /etc/sudoers.d/ipmdhcp
```

Note that you can change the password `admin` of the web service using the command below:

```
# htpasswd /usr/local/nessy2/www/php/cmd/dhcp/.htpasswd admin
```

If you have not installed the DNS package or are not planning on installing it, you must now follow the procedure below.

To complete the DHCP package installation on Debian if the DNS package is not installed

1. If relevant, open an SSH session.
2. Allow SNMP access to the DNS statistics: append the file `/etc/snmp/snmpd.conf` with the following line.

```
view systemonly included .1.3.6.1.4.1.2440
```

3. Start the SNMP daemon, using the following command:

```
# service snmpd start
```

4. Create a self-signed certificate for Apache, using the following commands:

```
# cd /etc/apache2  
# openssl genrsa -des3 -out server.key 4096  
# openssl req -new -key server.key -out server.csr  
# openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt  
# openssl rsa -in server.key -out server.key.insecure  
# mv server.key server.key.secure  
# mv server.key.insecure server.key
```

5. Activate the SSL mode in Apache using the following command:

```
# a2enmod ssl
```

6. Make sure that a symbolic link to the default VirtualHost SSL configuration file is located in the folder `sites-enabled/`. If not, use the following command.

- a. For Debian 7 and prior:

```
# ln -s /etc/apache2/sites-available/default-ssl /etc/apache2/sites-enabled/
```

- b. For Debian 8 and Debian 9:

```
# a2ensite default-ssl
```

7. Configure the web services.

- a. For Debian 7 and prior, in the file `/etc/apache2/sites-enabled/default-ssl`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

```
<VirtualHost *:443>

ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    AllowOverride All
</Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/apache2/server.crt
SSLCertificateKeyFile /etc/apache2/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

- b. For Debian 8 and Debian 9, in the file `/etc/apache2/sites-enabled/default-ssl.conf`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

```
<VirtualHost *:443>

ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    Require all granted
    AllowOverride Authnconfig
    Options Indexes FollowSymLinks
</Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/apache2/server.crt
SSLCertificateKeyFile /etc/apache2/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

8. Disable the default site in Debian Apache configuration using the following commands.

a. For Debian 7 and prior:

```
# cd /etc/apache2/sites-enabled
# unlink 000-default
```

b. For Debian 8 and Debian 9:

```
# a2dissite 000-default
```

9. Restart Apache using the following command:

```
# service apache2 restart
```

10. Make sure that the *ipmdhcp* package is running using the following command:

```
# service ipmdhcp status
```

If it is not running, use the following command:

```
# service ipmdhcp start
```

Once the configuration is complete, you can add an EfficientIP Package DHCP server to manage your ISC server from SOLIDserver GUI. Refer to the procedure in the section [Adding an ISC DHCP Server](#) for more details.

Installing the EfficientIP DHCP package for Linux CentOS/RedHat 7 and prior - 64 bits

You must take into account the [prerequisites](#) before [installing](#) a DHCP CentOS/RedHat package.

Prerequisites

- The DHCP package file, *ipmdhcpxx-y.y.y-redhatx.x86_64.rpm*, whose name provides you with a number of information separated by hyphens or a point: the type of package (*ipmdhcpxx*: a DHCP package with a DHCP in version *xx* where *xx* is *x dot x*), the version of SOLIDserver (*y.y.y*); the version of RedHat (*redhatx*) and finally the RedHat architecture (*x86_64*).

In the procedure below, this file is referred to as *<ipmdhcpxx-y.y.y-redhatx.x86_64.rpm>*.

- The EfficientIP ISC package platform must have at least 20 Mb of free disk space.
- The EfficientIP ISC package may need certain libraries of your operating system, you must have a *shell* access with *root* login in local, via *ssh*⁶ on the server to be installed.
- You must make sure that no other DNS/DHCP service on your Linux is running : it would interfere with the BIND/ISC package installation.
- You must make sure that SOLIDserver and RedHat/CentOS are set to the same time and date.
- You must make sure that Apache server is up-to-date.
- You must make sure that HTTPS (port 443), the DHCP service (port 67) and the failover ports (647-667 and 847-867) are not blocked by a network filtering process (firewall).

If your Apache configuration already uses the port 443, you have to create an additional IP-based VirtualHost dedicated to the DNS/DHCP management.

⁶You could also connect via *telnet* but, for security purposes, we recommend that you favor *ssh*.

Installing the EfficientIP DHCP Package

You can install the EfficientIP DHCP Package on both RedHat and CentOS Linux.

If you have not installed the DNS packages yet, you need to:

1. Follow the procedure [To install the EfficientIP DHCP Package on RedHat and CentOS](#).
2. Follow the procedure [To complete the DHCP package installation on RedHat/CentOS if the DNS package is not installed](#).

If you already installed the DNS packages, you only need to follow the procedure [To install the EfficientIP DHCP Package on RedHat and CentOS](#) below.

The installation procedure below includes the commands that make the web services configurable.

To install the EfficientIP DHCP Package on RedHat and CentOS

1. Open an SSH session.
2. Stop and disable your DHCP software, using the commands below.

- On RedHat:

```
# service isc-dhcp-server stop
# update-rc.d -f isc-dhcp-server remove
```

- On CentOS:

```
# service dhcpd stop
# chkconfig dhcpd off
```

3. Install the dependency packages, **ONLY** if you have not installed the EfficientIP DNS package, using the following commands:

```
# yum install mod_ssl php php-pdo sudo net-snmp sqlite
```

4. Install the EfficientIP DHCP package, using the following command:

```
# rpm -ivh <ipmdhcpxx-y.y.y-redhatx.x86_64.rpm>
```

5. Make the web services configurable: in the directory `/etc/sudoers.d`, create the file `ipmdhcp` containing the line below.

```
apache ALL = NOPASSWD: /usr/local/nessy2/script/install_dhcpd_conf.sh, \
                /usr/local/nessy2/script/install_dhcpd6_conf.sh
```

6. Set the users access rights as follows:

```
# chmod 440 /etc/sudoers.d/ipmdhcp
```

Note that you can change the password `admin` of the web service using the command below:

```
# htpasswd -c /usr/local/nessy2/www/php/cmd/dhcp/.htpasswd admin
```

If you have not installed the DNS package or are not planning on installing it, you must now follow the procedure below.

To complete the DHCP package installation on RedHat/CentOS if the DNS package is not installed

1. If relevant, open an SSH session.

2. Disable the firewall using the following commands.

- a. For RedHat/CentOS 6 and prior:

```
# service iptables stop
# chkconfig iptables off
```

- b. For RedHat/CentOS 7:

```
# service firewalld stop
# chkconfig firewalld off
```

3. Disable selinux. In the file `/etc/selinux/config`, modify the line `SELINUX=enforcing` to match the following one:

```
SELINUX=disabled
```

4. Reboot the system to take into account the selinux policy changes :

```
# reboot
```

5. In the file `/etc/sudoers`, disable `requiretty` by making it a comment as follows:

```
#Defaults requiretty
```

6. Allow SNMP access to the DHCP statistics. In the file `/etc/snmp/snmpd.conf`, in the section entitled Access Control, enter the lines:

```
master agentx
view systemview included .1.3.6.1.4.1.2440
#You may need to specify another view, AllView or a custom one,
#if you edited the default SNMP configuration.
```

7. Start the SNMP daemon, using the following command:

```
# service snmpd start
```

8. Create a self-signed certificate for Apache, using the following commands:

```
# cd /etc/httpd
# openssl genrsa -des3 -out server.key 4096
# openssl req -new -key server.key -out server.csr
# openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt
# openssl rsa -in server.key -out server.key.insecure
# mv server.key server.key.secure
# mv server.key.insecure server.key
```

9. Configure the web services. In the file `/etc/httpd/conf.d/ssl.conf`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

- a. For RedHat/CentOS 6 and prior:

```
<VirtualHost *:443>
ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    AllowOverride All
</Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
```

```
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/httpd/server.crt
SSLCertificateKeyFile /etc/httpd/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

b. For RedHat/CentOS 7:

```
<VirtualHost *:443>
ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    Require all granted
    AllowOverride Authconfig
    Options Indexes FollowSymLinks
</Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and
# must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/httpd/server.crt
SSLCertificateKeyFile /etc/httpd/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

10. Restart Apache using the following command line:

```
# service httpd start
```

11. Make sure that the *ipmdhcp* package is running using the following command line:

```
# service ipmdhcp status
```

If it is not running, use the following command line:

```
# service ipmdhcp start
```

Once the configuration is complete, you can add an EfficientIP Package DHCP server to manage your ISC server from SOLIDserver GUI. Refer to the procedure in the section [Adding an ISC DHCP Server](#) for more details.

Upgrading Packages

No matter the package version, to upgrade packages:

1. You must **uninstall your current packages**.
2. You must **install the new package** following the prerequisites and procedures detailed above in the section [Managing EfficientIP ISC Linux Packages](#).

Adding an ISC DHCP Server

Once you successfully installed your EfficientIP ISC Linux Package, you need to install an EfficientIP DHCP package and configure it according to your needs to manage BIND servers through the GUI.

The EfficientIP DHCP package is only available for DHCPv4 servers managed with SSL.

To add an ISC DHCP server for a Linux package

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the menu, select **+ Add > Server > EfficientIP DHCP Package**. The wizard **Manage a DHCP server** opens.
3. If you or your administrator created classes at the server level, in the list **DHCP server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the following fields to set up the basic server configuration:

Table 25.8. DHCP server basic parameters

Field	Description
DHCP server name	In this field, fill in a FQDN name for your server. This field is required.
Management IP address	In this field, fill in the IPv4 address of your server. This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DNS. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> . This field is optional.

5. If you modified the SSH login and password, tick the box **Configure enrollment parameters**. If not, go to step 6.

Once you ticked the box, the fields **Login** and **Password** appear. By default they both contain *admin*, edit them to make sure that the SSL credentials match your SSH credentials.

6. If you want to edit the server SNMP parameters⁷, tick the box **Configure SNMP monitoring parameters**. If not, go to step 7.

⁷The SNMP protocol parameters are used to monitor and retrieve the server statistics.

Once you ticked the box, the following fields appear:

Table 25.9. SNMP parameters used to monitor the server statistics

Field	Description
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
Use TCP	Tick the box if you want to use the TCP protocol instead of the UDP when the network link is not reliable.
SNMP profile	The SNMP profile used to retrieve the statistics. By default, <i>standard v2c</i> is selected. The list contains the default profiles (<i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i>) and the ones you may have created. Each profile has its own level of security and enables the definition of a global security policy. For more details, refer to the section Managing SNMP Profiles .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.

- Depending on the administrator configuration, you may be able to configure advanced properties according to the table below.

Table 25.10. Advanced properties configuration

Field	Description
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The server is listed.

Once the EfficientIP Package server is added, you can manage your ISC server in Linux from the GUI.

Editing a DHCP Server

To edit any kind of DHCP server configuration, you need to open its properties page and edit the relevant panel(s). Before proceeding, note that:

- The panels that do not contain the button *EDIT* cannot be edited.
- The basic server parameters are detailed in the addition procedures of each kind of physical server.
- For server managed via a smart architecture, some parameters can only be edited when you edit the smart. For instance, on EfficientIP servers managed via a smart, the box *Isolated* and the advanced properties parameters are only available when you edit the smart.

- The SNMP protocol is no longer supported as managing protocol for a server, so **editing it automatically changes the management protocol to use SSL instead of SNMP**. This operation is **non-reversible**.
- DHCP servers in version 7.0 are managed through a dedicated service account using a randomly generated password. Note that:
 - A server can be managed by only one appliance. To switch the appliance managing the server, you need to edit your sever and type in your credentials again.
 - For servers added in previous versions, after an upgrade to version 7.0, to switch to this new management system, you need to edit the servers.

If you want to add, edit or delete DHCP options, refer to the next section [Configuring DHCP Options at Server Level](#).

To edit a DHCP server

1. In the sidebar, go to  **DHCP > Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. At the end of the line of the server of your choice, click on . The properties page opens.
4. Open all the panels using .
5. In the panel of your choice, click on **EDIT**. The corresponding wizard opens.
6. Make the changes you need. Click on **NEXT** if need be to get to last page of the wizard.
7. For an EfficientIP server, in the panel Main Properties:
 - a. To manage a server that is already managed by another appliance, tick the box Configure enrollment parameters. The field "**Admin" account password** appears.
 - b. In the field "**Admin" account password**, enter your SSH password.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and refreshes.

Configuring DHCP Options at Server Level

You can configure DHCP options at server level, one server at a time. Keep in mind that:

- The DHCP options of a server are inherited by the scopes, groups, ranges, leases and/or statics it manages.
- On Microsoft DHCP servers, encapsulated DHCP options are not supported.
- You can aggregate range or static options on the scopes of a server.

For more details regarding the DHCP options configuration, refer to the chapter [Configuring DHCP Options](#) and/or to the appendix [DHCP Options](#).

Editing the DHCP Options of a Server

From the properties page of your EfficientIP servers, you can set DHCP options.

To edit the DHCP server options

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on . The properties pages opens.
3. In the upper right corner, click on .
4. In the panel **DHCP Options**, click on **EDIT**. The wizard **Configure DHCP options** opens.
5. In the drop-down list **Options category**, select the option type of your choice. The wizard refreshes.
6. Edit the option(s) of your choice.
7. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Aggregating DHCP Options from Ranges or Statics

If you imported an external DHCP configuration and/or a specific DHCP option is configured across all the ranges or statics of a scope, you can aggregate it, is apply it, on the scope managing the ranges or statics.

This aggregation automates a homogeneous configuration of DHCP options on each of the scopes of a specific server. Keep in mind that it:

1. Analyzes the DHCP options configured on the ranges or statics with IP address of a scope.
2. If one DHCP option is configured and has the same value on all the ranges or statics with IP address of a scope:
 - a. The DHCP option is applied to the parent scope.
 - b. The DHCP option is deleted at range or static level, as it is automatically propagated from the scope down.
3. If one DHCP option is configured on all the ranges or statics with IP address of a scope, but their value is not the same on all the objects, the option is not aggregated at scope level.

To aggregate range options

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Tick the server(s) of your choice.
3. In the menu, select  **Edit > Aggregate range options**. The wizard opens.
4. Click on **OK** to complete the operation. The report opens and closes.

To aggregate static options

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. Tick the server(s) of your choice.
3. In the menu, select  **Edit > Aggregate static options**. The wizard opens.
4. Click on **OK** to complete the operation. The report opens and closes.

Deleting a DHCP Server

At any time you can delete a server from the page *All servers*, this way you stop managing them from SOLIDserver.

Keep in mind that **you cannot delete a physical server if it is managed by a smart architecture.**

To delete a DHCP server

1. In the sidebar, go to  **DHCP > Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. Tick the server(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on  to complete the operation. The report opens and closes. The server might be marked  **Delayed delete** until it is no longer listed.

Defining a DHCP Server as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a server as one of the resources of a specific group allows the users of that group to manage the server in question as long as they have the corresponding rights granted.

Granting access to a server as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 26. Managing DHCP Scopes

SOLIDserver scopes constitute a level in the DHCP module and are used to determine the topology of the network, apply DHCP options for a routable domain, describe network clients, and indicate the addresses to be allocated to certain clients. In order to use the DHCP service, each terminal network to be served must have a DHCP scope that matches its IP address and its netmask (size). When a DHCP server serves clients with local physical networks, the scope is easily assimilated to its broadcast domain. A scope belongs to a DHCP server and can contain several DHCP ranges.

Browsing Scopes

The scopes are required in a dynamic addressing DHCP hierarchy, they are created directly within a DHCP server.

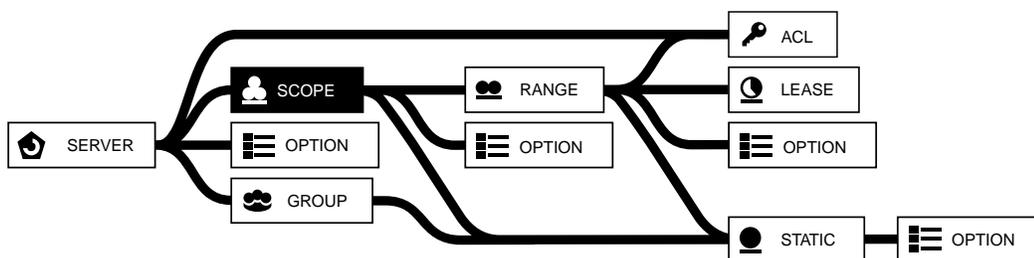


Figure 26.1. The scope in the DHCP hierarchy

Browsing the Scopes Database

To display the list of DHCP scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. You can filter the list to display a specific scope. Type in the data you are looking for in the search engine of the columns **Name** and/or **Address** and click on **REFRESH** to only display the scope matching your search.
4. To display the scopes of a server, in the column **Server**, click on the server of your choice. The page refreshes.

To display a DHCP scope properties page

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the scope of your choice, click on **ⓘ**. The properties pages opens.

Customizing the Display on the Page All Scopes

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that in IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the DHCP Scope Statuses

The column **Status** provides information regarding the scopes you manage.

Table 26.1. DHCP scope statuses

Status	Description
 OK	The scope is configured
 <i>Delayed create</i>	The creation or update is delayed due to a scheduled configuration of the server. The creation is automatically done after maximum of 1 minute.
 <i>Delayed delete</i>	The deletion is delayed due to a scheduled configuration of the server. The deletion is automatically done after maximum of 1 minute.

Adding a DHCP Scope

The addition of a scope to a DHCP server defines a new extension to the network's topology. Once created, the DHCP server is ready to receive a range of dynamic addresses.

Note that you can also import scopes, for more details refer to the section [Importing Scopes](#).

You can add as many scopes as you need on the page All scopes of all the server or a specific one. The procedure below creates a scope from the page All scopes with no server, otherwise the creation process can be slightly shorter.

To add a DHCP scope

1. In the sidebar, go to [DHCP > Scopes](#). The page **All scopes** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on  **Add**. The **Add a DHCP scope** or **Add a DHCPV6 scope** wizard opens.
4. In the list **DHCP server**, select the DHCP server in which you want to create the scope.
5. Click on . The next page of the wizard appears.
6. If you or your administrator created classes at the scope level, in the list **DHCP scope class** select a class or *None*. Click on . The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. Click on . The next page of the wizard appears.
8. Fill in the following fields to configure the scope parameters:

Table 26.2. DHCP scope configuration parameters

Field	Description
Name	Name the scope. This field is required.
Network address	Type in the scope address. This field is required.
Netmask	This field is only available for V4. Select the shared physical network. This field is required.
Prefix	Select the scope prefix. By default, the prefix is selected depending on the netmask you chose. The prefix selected edits the value of the field <i>Netmask</i> . This field is required.
Shared network	Type in the first digits of the address of an existing shared network, the auto-completion provides a list matching your search. Select the shared network of your choice. If you leave the field empty, a new shared network is created and named after the scope <i>start_address/prefix</i> .
Failover	Select a failover channel. By default, <i>None</i> is selected. This field is optional and only available when you add a scope to a smart architecture, other than a <i>Single-Server</i> .
DHCP scope space name	Select one of your existing IPAM spaces: the scope is part of it. This field is optional.

- Depending on the administrator configuration, you may be able to configure advanced properties according to the table below.

Table 26.3. Advanced properties configuration

Field	Description
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The scope is listed.

Editing a DHCP Scope

To edit the main properties of a DHCP scope

- In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
- On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
- At the end of the line of the scope of your choice, click on . The properties page opens.
- In the panel **Main properties**, click on . The **Edit a DHCP scope** wizard opens.
- If you or your administrator created classes at the scope level, in the list **DHCP scope class** select a class or *None*. Click on . The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. You can modify the **Name**, **Shared network**, **DHCP scope space name** and **Advanced Properties**.
7. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Configuring DHCP Options at Scope Level

You can configure DHCP options at scope level, individually or in bulk. Keep in mind that:

- The DHCP options of a scope may be inherited from its server.
- The options set at scope level override the options set on at server level.
- All the DHCP options you configure at scope level are inherited by all the ranges, leases and/or statics it manages.
- You can only edit DHCP options of several scopes at once on DHCPv4 servers.

For more details regarding the DHCP options configuration, refer to the chapter [Configuring DHCP Options](#) and/or to the appendix [DHCP Options](#).

Editing the DHCP Options of a Scope

From the properties page of your scopes, you can set DHCP options.

To edit a scope DHCP option

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the scope of your choice, click on **⚙**. The properties pages opens.
4. In the upper right corner, click on **☰**.
5. In the panel **DHCP options**, click on **EDIT**. The **Configure DHCP options** wizard opens.
6. In the drop-down list **Options category**, select the option type of your choice. The wizard refreshes.
7. Edit the option(s) of your choice.
8. Click on **OK** to complete the operation. The report and closes. The changes are listed in the panel.

Performing Option Changes on Several Scopes At Once

From the page **All scopes**, you can be set, replace or delete DHCP options on all the DHCPv4 scopes you select at once.

To add a DHCP option to one or several scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the scope(s) of your choice.

5. In the menu, select **↗ Edit > Option > Add**. The wizard **Add DHCP scope options** opens.
6. In the drop-down list **Option name**, select an option.
7. In the field **Value**, type in the relevant value.
8. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the scope properties page, the panel **DHCP options** lists the new DHCP option and its value.

To edit the value of a DHCP option on one or several scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the scope(s) of your choice.
5. In the menu, select **↗ Edit > Option > Replace**. The wizard **Replace DHCP scope options** opens.
6. In the drop-down list **Option name**, select the option which value you want to replace.
7. In the field **Replace**, specify the value you want to change.
8. In the field **By**, type in the new option value.
9. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the scope properties page, the panel **DHCP options** displays the new value of the DHCP option.

To remove a DHCP option from one or several scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the scope(s) of your choice.
5. In the menu, select **↗ Edit > Option > Delete**. The wizard **Delete DHCP scope options** opens.
6. In the drop-down list **Option name**, select an option.
7. In the field **Option value filter**, type in the option value.
8. Click on **OK** to complete the operation. The report opens and closes. On the scope properties page, the panel **DHCP options** no longer displays the DHCP option.

Defining a Specific IPAM Space for a Scope

At any time you can associate DHCP scopes with IPAM spaces, either to define a space if you did not when you added the scope or to set a different one.

Defining a specific space at scope level allows to apply policy rules from the IPAM module to several addresses and avoid any overlapping of ranges and spreads of reserved addresses.

You can associate scopes with a space one scope at a time or several scopes at once.

To edit a scope and change the space it is associated with

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Right-click on the scope of your choice. The contextual menu opens.
4. Click on **⚙**. The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The **Edit a DHCP scope** wizard opens.
6. In the drop-down list **DHCP scope space name**, select the space of your choice or *None* to remove the association.
7. Click on **OK** to complete the operation. The report opens and closes. The panel displays the new space name.

To set the space of one or several scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the scope(s) you want to edit. The scope(s) can already be associated with a space.
4. In the menu, select **⚙ Edit > Set > Space**. The wizard **Edit the scope space** opens.
5. In the drop-down list **Space**, select the space you want to associate with your scope(s) or *None* to remove the association.
6. Click on **OK** to complete the operation. The report opens and closes. The new space is listed in the column **Scope space**.

Once a scope is associated with a space, you can execute the DHCP to IPAM replication and associate scopes with existing networks or create the corresponding network. As you can create several terminal networks managing the same IP addresses in separate spaces, you can associate these networks with scopes belonging to distinct DHCP servers.

Defining a Specific Failover Channel for a Scope

At any time you can associate DHCPv4 scopes with a specific failover if they belong to a smart architecture, other than a *Single-Server*. For more details regarding failover channels, refer to the chapter [Managing Failover Channels](#).

You can associate scopes with a failover channel one scope at a time or several scopes at once.

To edit a scope and change the failover it is associated with

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** depending on your needs. The page refreshes and the button turns black.
3. Right-click on the scope of your choice. The contextual menu opens.
4. Click on **⚙**. The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The **Edit a DHCP scope** wizard opens.

6. In the drop-down list **Failover**, select the failover of your choice or *None* to remove the association.
7. Click on **OK** to complete the operation. The report opens and closes. The panel displays the new space name.

To set the failover of one or several scopes

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the scope(s) you want to edit. The scope(s) can already be associated with a space.
4. In the menu, select **Edit > Set > Failover channel**. The wizard **Configure DHCP scopes failover** opens.
5. In the drop-down list **Failover**, select the failover channel you want to associate with the scope(s) or *None* to remove the association.
6. Click on **OK** to complete the operation. The report opens and closes. The new space is listed in the column **Scope space**.

Replicating Scope Data in the IPAM

The option *IPAM replication* allows to decide in which space you want to replicate scope data. If you want to associate a scope with a space or change the space it is associated with, refer to the section [Defining a Specific IPAM Space for a Scope](#).

Before replicating your DHCP data into the IPAM, keep in mind that:

- The option *IPAM replication* behaves as follows:
 - If you specify a space that has no matching network, a network is created, it is named like the scope.
 - If you do not specify any space, with the value *None*, SOLIDserver automatically creates a network matching the scope in the first space that can receive it.
- The option *IPAM replication* can only specify a target space for scopes that are not configured with any space yet.

Either no space was specified when you added it or you set the space of a scope and selected *None*, as detailed in the procedure [To set the space of one or several scopes](#).

- The option *IPAM replication* does not change the associated space, it sets a replication association. If you want to change the target of a scope already associated with a space, you must edit as detailed in the section [Defining a Specific IPAM Space for a Scope](#).

To replicate scope data in the IPAM

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the scope(s) of your choice.
4. In the menu, select **Edit > IPAM replication**. The wizard **IPAM replication** opens.

5. In the drop-down list **Target space**, select a space or *None*. For more details, refer to the [option behavior details](#).
6. Click on **OK** to complete the operation. The report opens and closes. The new space is listed in the column **Scope space**.

Note that you can also replicate DHCP range and static data to the IPAM. For more details, refer to the sections [Replicating Range Data in the IPAM](#) and [Replicating Static Data in the IPAM](#).

Configuring Multiple Scopes for a Network Segment

You can make several scopes serve an entire network segment as a single entity with shared networks.

For instance, if you configured multinetting on your network and one DHCP server answers client requests on a single physical network that has multiple IP networks in use. With a shared network containing several scopes, the server identifies that a client request was sent from one of the scopes and that it has many available IP addresses to choose from and assign to the client.

If dynamic DHCP ranges appear within scopes using the same shared network, all address ranges are offered independently. Once the first range is full, the ranges that are declared within the same shared network are used one after the other until all addresses are used.

To add a DHCP scope to a shared network

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. At the end of the line of the scope of your choice, click on **⊞**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The **Edit a DHCP scope** wizard opens.
4. If you created classes at scope level, in the list **DHCP scope Class**, select a class if needed.
5. Click on **NEXT**. The last page of the wizard appears.
6. In the field **Shared network**, if you did not set a shared network when creating the scope, the start IP address and the prefix of the scope you are editing is displayed. Type in the first digits of the address of an existing shared network, the auto-completion provides a list matching your search. Select the shared network of your choice.
7. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Copying or Moving DHCPv4 Scopes

To assist you in the management, scopes can be copied and moved from one server to the other. In both cases, make sure that the IP addresses they manage is not already managed by another scope in the target space.

Migrating a scope also migrates the DHCP ranges and statics with IP address it contains. As for the statics without IP migration, refer to the section [Copying a DHCPv4 Static Without IP](#).

Keep in mind that if your physical server is managed via a smart, only the scope created on the smart can be duplicated or moved.

To copy a scope to another server

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Tick the scope(s) you want to duplicate.
4. In the menu, select **Edit > Migrate**. The **Copying/Moving scopes** wizard opens.
5. In the drop-down list **Method**, select *Copy*.
6. In the drop-down list **Target server**, select the server or smart architecture of your choice.
7. Click on **OK** to complete the operation. The report opens and closes. The page **All scopes** is visible again. Both scopes are listed and they share the same name, start address and end address. The duplicate scope is in *Delayed create* in the target server.

To move a scope to another server

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. Tick the scope(s) you want to duplicate.
4. In the menu, select **Edit > Migrate**. The **Copying/Moving scopes** wizard opens.
5. In the drop-down list **Method**, select *Move*.
6. In the drop-down list **Target server**, select the server or smart architecture of your choice.
7. Click on **OK** to complete the operation. The report opens and closes. The page **All scopes** is visible again. The scope is no longer listed as part of the first server. It now belongs to the selected target server.

Deleting a DHCP Scope

The deletion of a DHCP scope makes the address ranges and leases it contains disappear. This deletion restricts the network's topology.

To delete a DHCP scope

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be. For more details, refer to the procedure [To display the list of DHCP scopes](#).
4. Tick the scope(s) you want to delete.
5. In the menu, click on **Delete**. The Delete wizard opens.
6. Click on **OK** to complete the operation. The report opens and closes. The scope is no longer listed.

DHCP Relay Agents

Rather than directly connecting the DHCP server to every network segment it serves, it is possible to configure a DHCP relay agent on each network segment. Relay agents are configured with a list of one or more DHCP servers, two servers must be configured for the DHCP failover. When a relay agent receives a message from a DHCP client on a particular network segment, it records the IP address of the interface on which it received the request in the field *GiAddr* of the message, and then it forwards the message to the DHCP server. From there, the server directly responds to the client.

The DHCP relay is a mechanism that allows the transfer of DHCP/BOOTP messages between clients and servers of different networks. The routers used to interconnect these networks possess for the most part the functionality of TCP/IP relay agents. To conform to RFC 1542 and deal with the relay agent, each router must be able to recognizing BOOTP and DHCP messages and relaying them in an appropriate manner. A router equipped with the capacities of a BOOTP relay agent generally relays DHCP packets, as well as all BOOTP packets transmitted on the network. SOLIDserver supports DHCP relay transparently. If a scope has the same network address as one of the interfaces of the DHCP server, then it is a local scope. This means that it belongs to the same broadcast domain as the DHCP server. Otherwise, it is a relay scope.

BOOTP / DHCP relay on Cisco devices (IP helper)

If we consider two DHCP servers, with one on the network *191.24.1.0* and the other one on *110.44.0.0*. To allow the IP broadcast from all hosts to be forwarded in unicast toward both servers, set the configuration below.

```
interface ethernet1
 ip helper-address 191.24.1.45
 ip helper-address 110.44.0.125
```

BOOTP / DHCP relay on Juniper devices (IP helper)

If we consider that a DHCP server is in VLAN *20* with the IP address *20.20.20.2*, the DHCP client's computer is in VLAN *10* and a Juniper switch is configured as DHCP relay and performs inter VLAN routing between the VLANs *10* and *20*, set the configuration below.

```
set vlans vlan10 vlan-id 10 set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members
vlan10 set vlans vlan10 13-interface vlan.10 set interfaces vlan unit 10 family inet address
10.10.10.1/24 set vlans vlan20 vlan-id 20 set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members vlan20 set interfaces vlan unit 20 family inet address 20.20.20.1/24 set vlans vlan20
13-interface vlan.20 set forwarding-options helpers bootp server 20.20.20.2 set forwarding-options
helpers bootp interface vlan.10
```

BOOTP / DHCP relay on HP devices (IP helper)

If we consider a DHCP server with the IP address *10.10.20.3* IP and a DHCP client's computer is in VLAN *40*, set the configuration below.

```
vlan 40 ip helper-address 10.10.20.3
```

Defining a DHCP Scope as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a scope as one of the resources of a specific group allows the users of that group to manage the scope in question as long as they have the corresponding rights granted.

Granting access to a scope as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 27. Managing Fixed Reservations

DHCP fixed reservation relies on [statics](#) and [groups](#).

Managing DHCP Statics

A DHCP static is essentially a lease reservation that ensures a specified client always uses the same IP address on a network. For clients who require a constant IP address, you can manually configure an IP address or assign a DHCP static reservation, reserving a static allows to take advantage of the DHCP options.

Static reservations can match DHCP, PXE or BOOTP clients based on based on their MAC address or *DHCP-client-identifier*. These reservations can belong to a DHCP server directly, a group, a scope or a range. All the DHCP options set on these containers apply to the reservations, so if you edit the DHCP options, the devices configured with the options are automatically updated when they request the lease renewal.

Every DHCP static reservation must have a unique name which is usually used to identify it but, in particular contexts, can be used to enforce the client's hostname.

When it comes to statics, there is a main difference between DHCP managing IPv4 and IPv6 addresses. DHCPv6 introduces a new piece of information, the DHCP Unique Identifier (DUID). It should not exceed 128 bits in total and allows to identify a client rather than an equipment. It contains the MAC address, therefore this address is not a unique independent set of numbers anymore, it corresponds to the last 48 to 64 bits of the DUID depending on its type.

There are three different types of DUID:

- DUID based on Link Layer (LL).
- DUID based on Link-Layer Address Plus Time (LLT).
- DUID Assigned by Vendor Based on Enterprise Number (EN).

The DUID default structure goes like this:

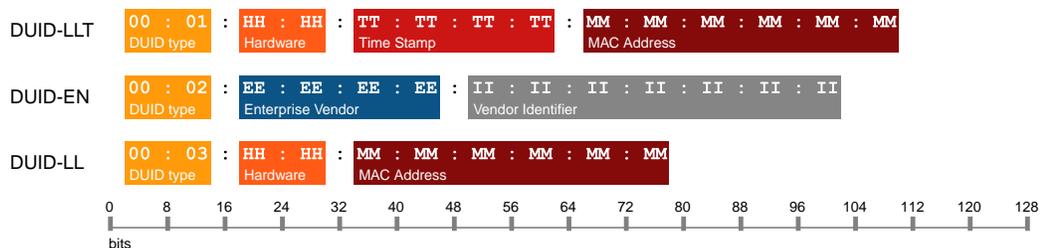


Figure 27.1. Three different structures of DUID

DHCP servers in IPv4 use the MAC address specified during the static reservation to identify the clients' IP address and allocate them a lease as well as soon as they are visible on the network. Once the lease is allocated, if the IPAM and DNS replication are configured, the data is sent to the IPAM and the DNS.

Browsing DHCP Statics

The DHCP statics are the end point of a DHCP fixed reservation strategy. They can belong to a DHCP group or directly to a DHCP server.

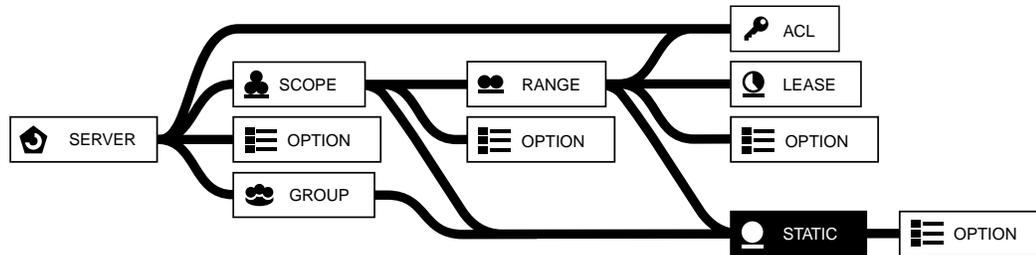


Figure 27.2. The static reservation in the DHCP hierarchy

Browsing the DHCP Statics Database

To display the list of statics

1. In the sidebar, go to [DHCP > Statics](#). The page **All statics** opens.
2. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.
3. To display the list of statics of a DHCP server or smart architecture, in the column **Server**, click on the name of your choice. The page refreshes.

To display a DHCP static properties page

1. In the sidebar, go to [DHCP > Statics](#). The page **All statics** opens.
2. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the static of your choice, click on [\[icon\]](#). The properties page opens.

On the statics properties page you can find the following information in separate panels:

- **Main properties:** sums up all the information filled in during the static creation (DHCP server, scope and group, its name, IP address, client DUID or MAC address, class).
- **Audit:** displays all of the changes performed on the static by the user logged in. If they belong to a group with access to the changes from all users, the panel displays all the operations ever performed. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).
- **DHCP options:** displays all the DHCP options you can define. None of the default options are listed except for the type of DHCP server. For more details, refer to the chapter [Configuring DHCP Options](#).

Customizing the Display on the Page All Statics

Users of the group *admin* can create customized column layouts. The button [\[icon\] Listing template](#), on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

As EfficientIP DHCP servers manage IPv4 static reservations like leases, two new columns were added to the page *All statics*: **Last seen** that indicates the last time the client was connected and **Expiration** that indicates when the lease expires.

Keep in mind that in IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the DHCP Static Statuses

The column **Status** provides information regarding the statics you manage.

Table 27.1. DHCP static statuses

Status	Description
✔ OK	The static is operational.
⌚ <i>Delayed create</i>	The creation is delayed due to a scheduled configuration of the server. The creation is automatically done after maximum of 1 minute.
⌚ <i>Delayed delete</i>	The deletion is delayed due to a scheduled configuration of the server. The deletion is automatically done after maximum of 1 minute.

Adding a DHCP Static

Whether you work with DHCPv4 or DHCPv6, the statics reservation process is the same: it identifies clients through their MAC to grant them permanent access to a server and a set of DHCP options inherited from their container. You can associate the client's MAC address with a specific IP address on your network or associate it with a static without IP address, this grants your client access to one server or to the servers connected to a failover channel.

During the static reservation you must specify a MAC address and choose its type, the type modifies the addresses display on the page. For more details regarding the supported MAC addresses types, refer to the appendix [MAC Address Types References](#).

Keep in mind that:

- If you or your administrator configured IPAM to DHCP advanced properties, new statics may be added for every new IP address created in the IPAM. For more details, refer to the chapter in the section IPAM Advanced properties of the chapter [Managing Advanced Properties](#).
- The rule *022 Check DHCP static duplicate hostnames* allows to automatically check that two different DHCPv4 statics do not have the same name on one DHCP server. This rule is enabled by default and triggered when you are about to add a new static.

Note that you can also import statics, for more details refer to the section [Importing Statics](#).

Adding DHCPv4 Statics

To set up static reservation you need a user identifier, the equipment MAC address, that you can associate with an IP address.

EfficientIP DHCP servers manage the statics with IP address like leases:

- **Prerequisites:**
 - EfficientIP DHCP servers in version 6.0.0 or more. One on its own or several servers via a smart architecture, the failover cannot support the management of statics like leases if one of the servers is in a version prior to 6.0.0.

- **New behavior:**
 - Adding a static with IP address also creates the corresponding lease whenever the client is active on the network.
 - You can associate one MAC address with several static reservations as long as you set them each with a different IP address.
 - Once the lease is active, if you set the IPAM and DNS replication, the corresponding entries in the DNS are created. When the lease expires, the DNS entries are removed.
 - If the IP address of the static is replicated in the IPAM, it has a specific *Type* and *Status* on the page All addresses. For more details, refer to the section [Understanding the IP Address Type and Status](#).
- **Limitations:**
 - The new static reservation management cannot be set on Microsoft DHCP servers, Nominum DCS servers or EfficientIP server in versions prior to 6.0.0.
 - If you migrated from 5.0.4, the statics configured with the option *host-name* are renamed: the value of the option is used as the static name. If the static were not configured with the option, their name is used as value of the option *host-name*.

To add a DHCPv4 static

1. In the sidebar, go to [DHCP > Statics](#). The page **All statics** opens.
2. On the right-end side of the menu, click on [\[v4\]](#). The page refreshes and the button turns black.
3. In the menu, click on [+ Add](#). The wizard **Add a DHCP static** opens.
4. In the list **DHCP server**, select the DHCP server or smart architecture of your choice.
5. Click on [\[NEXT\]](#). The next page opens.
6. In the list **DHCP scope**, select a scope or *None* if you do not want to assign an IP address to your static, in this case the server DHCP options apply to the static.
7. Click on [\[NEXT\]](#). The next page opens.
8. If you or your administrator created classes at the static level, in the list **DHCP static class**, select a class or *None*. Click on [\[NEXT\]](#). The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. Configure the static using the table below:

Table 27.2. DHCPv4 static configuration parameters

Field	Description
Name	You can name the static reservation, this name is used as value of the DHCP option "host-name". This field is optional.
IP address	Type in the IPv4 address you want to assign to the device. This field is only displayed if you selected a scope.
MAC address	Type in the MAC address of the device that should use the DHCP static. This field is required.
MAC address type	Select the protocol associated with the MAC address. The protocol corresponding reference is displayed before the MAC address in the column MAC address . <i>Ethernet</i>

Field	Description
	is selected by default. If you select <i>Unknown</i> , the field <i>Type reference</i> appears. This field is required.
Type reference	Specify the reference number of the unknown MAC address type. If the reference you type in is already part of the database, it is automatically retrieved and visible when editing the static. This field is required.
Group name	You can select an existing DHCP group for the static. If you select a group, its DHCP options apply to the static. The list only contains <i>None</i> if you have not created any group yet. You can edit the DHCP group of one or several statics from the page <i>All statics</i> , for more details refer to the section Editing the DHCP Group of DHCPv4 Statics .

10. Click on **OK** to complete the operation. The report opens and closes. The static is listed.

If you added the static to an EfficientIP DHCP server with an IP address belonging to a range, once the client is connected they are listed on the page *All leases* and their information updates the DNS if the IPAM and DNS replication is set.

Adding DHCPv6 Statics

To set up a static reservation with DHCPv6, you need an IP address and a user identifier: the client DUID. Considering that the DUID can be quite long - for more details, refer to the introduction of the section [Managing DHCP Statics](#) - you have the possibility to either put it in full in the DUID field or put only the DHCPv4 equivalent of the MAC address, that is to say the last 48 to 64 bits, looking like *xx : xx : xx : xx : xx : xx* or *xx : xx : xx : xx : xx : xx : xx*.

To add a DHCPv6 static

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard **Add a DHCPV6 static** opens.
4. In the drop-down list **DHCP server**, select the DHCP server or smart architecture of your choice.
5. Click on **NEXT**. The next page opens.
6. In the list **DHCP scope**, select a scope or *None* if you do not want to assign an IP address to your static, in this case the server DHCP options apply to the static.
7. Click on **NEXT**. The next page the wizard opens.
8. If you or your administrator created classes at the static level, in the list **DHCP static class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. Configure the static using the table below:

Table 27.3. DHCPv6 static configuration parameters

Field	Description
Name	You can name the static reservation. This field is optional.

Field	Description
IP address	Type in the IPv6 address that the device should use. This field is only displayed if you selected a scope.
Client DUID	Type in the equipment DUID. If you do not specify a <i>Client DUID</i> , you must fill in the field <i>MAC address</i> .
MAC address	Type in the MAC address, that is to say the six sets of hexadecimal digits of the equipment DUID. If you do not specify the <i>MAC address</i> , you must fill in the field <i>Client DUID</i> .
MAC address type	Select the protocol associated with the MAC address. The protocol reference is displayed before the MAC address in the default column MAC address . <i>Ethernet</i> is selected by default. This field is required.
Group name	You can select an existing DHCPv6 group for the static. If you select a group, its DHCP options apply to the static. The list only contains <i>None</i> if you have not created any group yet.

- Click on **[OK]** to complete the operation. The report opens and closes. The static is listed.

Editing a DHCP Static

You can edit the properties of existing DHCPv4 or DHCPv6 statics from their properties page or put them in a DHCP group from the page *All statics*.

Editing the Properties of a Static

At any time, you can edit your statics from the contextual menu or their properties page.

To edit a DHCPv4 static

- In the sidebar, go to **🔍 DHCP > Statics**. The page **All statics** opens.
- Filter the list if need be.
- At the end of the line of the static you chose, click on **⚙️**. The properties page opens.
- In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DHCP static** opens.
- If you or your administrator created classes at the static level, in the list **DHCP static class**, select a class or *None*. Click on **[NEXT]**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Edit the information of your choice in the fields **Name**, **MAC address**, **MAC address type** and/or **Group name**.
- Click on **[OK]** to complete the operation. The report open and closes. The modifications are visible in the panel.

To edit a DHCPv6 static

- In the sidebar, go to **🔍 DHCP > Statics**. The page **All statics** opens.
- Filter the list if need be.
- At the end of the line of the static you chose, click on **⚙️**. The properties page opens.
- In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DHCP static** opens.

5. If you or your administrator created classes at the static level, in the list **DHCP static class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Edit the information of your choice in the fields **Name**, **IP address**, **Client DUID**, **MAC address**, **MAC address type** and/or **Group name**.
7. Click on **OK** to complete the operation. The report opens and closes. The modifications are visible in the panel.

Editing the DHCP Group of DHCPv4 Statics

At any time from the page *All statics*, you can put one or several DHCPv4 statics in the DHCP group of your choice, or edit the group they belong to.

To put DHCPv4 statics in a DHCP group

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. Filter the list if need be.
3. Tick the static(s) you want to put in a different group.
4. In the menu, select **Edit > Modify > Group**. The wizard **Modify the DHCP group of a static** opens.
5. In the drop-down list **DHCP group**, select the group of your choice.
6. Click on **OK** to complete the operation. The report opens and closes. The group is listed in the column **Group** of the static.

Replicating Static Data in the IPAM

At static level, the option *IPAM replication* allows to replicate static data in IPAM addresses. Before replicating your DHCP data in the IPAM, keep in mind that:

- You can only replicate statics with IP in the IPAM.
- The option *IPAM replication* behaves as follows:
 - The replicated DHCP information overwrites the IPAM information.
 - If the static belongs to a scope for which the IPAM replication has been set, its data is replicated as an IP address that belongs to the same terminal network.

Depending on your IPAM configuration, this IP address can belong to a pool, even if the pool is read-only.

- If the static belongs to a scope for which no replication has been set, its data is replicated as an IP address in the first terminal network or pool that can receive it. The pool can be read-only.
- If you replicate a static with IP that matches an existing IP address:
 - The IP address is renamed, it takes the name of the static.
 - The MAC address of the IP address is edited, it takes the MAC address of the static.
- If you replicate a static with IP that has no corresponding IP address, an IP address is created. It has the same name and MAC address than the static with IP.

- All the statics replicated edit the IP addresses Type to *DHCP static* and Status to *Reserved*.
- The option *IPAM replication* is independent at static level, replicating scopes or ranges does not automatically replicates to the statics they contain. You must select statics and execute the option to update the IPAM IP addresses with range data.

For more details regarding scope and range replication, refer to the sections [Defining a Specific IPAM Space for a DHCPv4 Scope](#) and [Replicating Range Data in the IPAM](#).

To replicate range data in the IPAM

1. In the sidebar, go to  **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Tick the range(s) of your choice.
4. In the menu, select  **Edit > IPAM replication**. The wizard **IPAM replication** opens.
5. Click on  to complete the operation. The report opens and closes.

The matching pool is visible on the IPAM page **All pools**.

Configuring DHCP Options at Static Level

You can configure DHCP options at static level, individually or in bulk. Keep in mind that:

- You can only edit DHCP options at static level on DHCPv4 servers.
- The DHCP options of a static may be inherited from its server, group, scope or range.
- The options set at static level override the options set on its container.
- By default, statics created on an EfficientIP DHCP server are configured with the option *host-name*, its value is the static name.
- If your statics belong to a DHCP group, editing the DHCP options of the group also edits the options of the statics. For more details, refer to the chapter [Managing DHCP Groups](#).

For more details regarding the DHCP options configuration, refer to the chapter [Configuring DHCP Options](#) and/or to the appendix [DHCP Options](#).

Editing the DHCP Options of a Static

From the properties page of your statics, you can set DHCP options.

To edit the DHCP options of a static

1. In the sidebar, go to  **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. At the end of the line of the static of your choice, click on . The properties page opens.
4. In the upper right corner, click on .
5. In the panel **DHCP options**, click on . The wizard **Configure DHCP options** opens.
6. In the drop-down list **Options category**, select the option type of your choice. The wizard refreshes.

7. Edit the option(s) of your choice.
8. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Performing Option Changes on Several Statics At Once

From the page All statics, you can set, replace or delete DHCP options on all the statics you select at once.

To add a DHCP option to one or several statics

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the static(s) of your choice.
5. In the menu, select **Edit > Option > Add**. The wizard **Add DHCP options to statics** opens.
6. In the drop-down list **Option name**, select an option.
7. In the field **Value**, type in its value.
8. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the static properties page, the panel **DHCP options** lists the new DHCP option and its value.

To edit the value of a DHCP option on one or several statics

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the static(s) of your choice.
5. In the menu, select **Edit > Option > Replace**. The wizard **Replace DHCP options of statics** opens.
6. In the drop-down list **Option name**, select the option which value you want to replace.
7. In the field **Replace**, specify the value you want to change.
8. In the field **By**, type in the new option value.
9. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the static properties page, the panel **DHCP options** displays the new value of the DHCP option.

To remove a DHCP option from one or several statics

1. In the sidebar, go to **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the static(s) of your choice.

5. In the menu, select  **Edit > Option > Delete**. The wizard **Delete DHCP options from statics** opens.
6. In the drop-down list **Option name**, select the option you want to delete.
7. In the field **Option value filter**, specify the option value.
8. Click on  to complete the operation. The report opens and closes. On the static properties page, the panel **DHCP options** no longer displays the DHCP option.

Copying a DHCPv4 Static Without IP

You can copy statics without IP from one server to the other, even if its MAC address is already declared in the target server. Keep in mind that if your physical server is managed via a smart, only the static created on the smart can be duplicated.

Statics with IP address are copied or moved when you migrate the scope they belong to.

To copy a static without IP in another server

1. In the sidebar, go to  **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. Filter the list of need be.
4. Tick the static(s) without IP you want to duplicate.
5. In the menu, select  **Edit > Migrate**. The wizard **Copying statics** opens.
6. In the drop-down list **Target server**, select the server or smart architecture of your choice.
7. Click on  to complete the operation. The report opens and closes. The page **All statics** is visible again. The static is displayed twice, in two different servers.

Deleting a DHCP Static

At any point, you can delete a static reservations. Keep in mind that:

- If you or your administrator configured IPAM to DHCP advanced properties, deleting IP addresses in the IPAM also deletes the corresponding DHCP statics. For more details, refer to the chapter in the section IPAM Advanced properties of the chapter [Managing Advanced Properties](#).
- If you or your administrator configured IPAM to DHCP and IPAM to DNS advanced properties, adding a static with IP adds a lease on the page *All leases* when the client is active, once the IPAM database is updated it updates the DNS to add the corresponding entries. **When you delete a static with IP, the corresponding DNS entries are deleted as well once the lease expires.**

To delete a DHCP static

1. In the sidebar, go to  **DHCP > Statics**. The page **All statics** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the static(s) you want to delete.

5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on to complete the operation. The report opens and closes. The static is no longer listed.

Managing DHCP Groups

The DHCP group allows to apply one or more parameters to a group of static declarations. Configuring a group is optional, the statics would work properly without separately as well. For clients with statically assigned addresses, or for installations where only known clients should be served, each such client must have a DHCP static declaration. If parameters are to be applied to a group of declarations which are not related strictly on a per-network basis, the group declaration can be used. Some sites may have departments which have clients on more than one IP network, but it may be desirable to offer those clients a uniform set of parameters which are different than what would be offered to clients from other departments on the same IP network. For clients that should be declared explicitly with DHCP static declarations, these declarations can belong to a DHCP group declaration along with the parameters which are common to that department.

You can add as many groups as you want but you cannot edit them. You can delete them and replace them with new ones.

You can only add DHCP groups on EfficientIP DHCP servers. If you display the content of any other type of DHCP server, the breadcrumb no longer displays the page All groups.

Browsing DHCP Groups

The DHCP groups are an optional level of the DHCP fixed reservation organization. They can be used to manage the statics DHCP options.

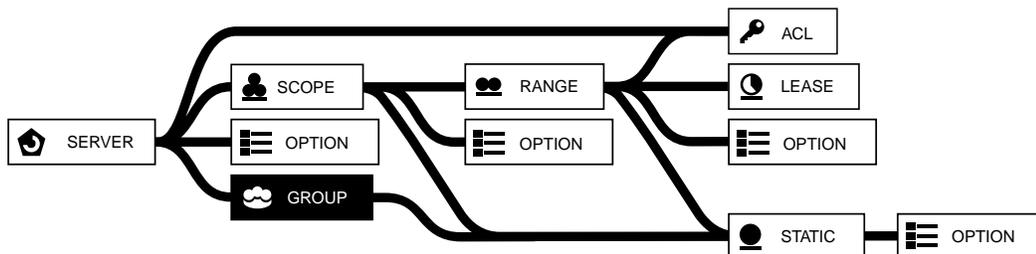


Figure 27.3. The group in the DHCP hierarchy

Browsing the Groups Database

To display the list of DHCP groups

1. In the sidebar, go to [DHCP > Groups](#). The page **All groups** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.

To display the list of groups of a DHCP server or smart architecture

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.

2. At the end of the line of the server or smart architecture of your choice, click on . The properties page opens.
3. In the breadcrumb, click on **All groups**. The page **All groups** of that server or smart architecture opens.

To display a DHCP group properties page

1. In the sidebar, go to  **DHCP > Groups**. The page **All groups** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the group of your choice, click on . The properties page opens.

Customizing the Display on the Page All Groups

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the DHCP Group Statuses

The column **Status** provides information regarding the groups you manage.

Table 27.4. DHCP group statuses

Status	Description
 <i>OK</i>	The group is operational.
 <i>Delayed create</i>	The creation is delayed due to a scheduled configuration of the server. The creation is automatically done after maximum of 1 minute.
 <i>Delayed delete</i>	The deletion is delayed due to a scheduled configuration of the server. The deletion is automatically done after maximum of 1 minute.

Adding a DHCP Group

At any point, you can add a group to an EfficientIP DHCP server. The purpose of a DHCP group is to apply a set of DHCP options to the statics that you want to manage with it.

Even if you already created statics, you have the possibility to put them in the group of your choice. For more details, refer to the section [Editing the DHCP Group of DHCPv4 Statics](#).

You can add a group from the page *All groups* with both DHCPv4 and DHCPv6.

To add a DHCP group

1. In the sidebar, go to  **DHCP > Groups**. The page **All groups** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. In the menu, click on  **Add**. The wizard **Add a DHCPV6 group** opens.
4. In the list **DHCP server**, select the DHCP server in which you want to add the group.
5. Click on . The next page of the wizard opens.

6. If you or your administrator created classes at the group level, in the list **DHCP group class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **DHCP group name**, type in the group name.
8. Click on **OK** to complete the operation. The report opens and closes. The group is listed.

Keep in mind that you cannot edit a group, you have to delete it.

Configuring DHCP Options at Group Level

At group level, you can configure DHCP options. They are inherited by statics it manages.

For more details regarding DHCP options, refer to the chapter [Configuring DHCP Options](#) and/or to the appendix [DHCP Options](#).

To edit the DHCP options of a group

1. In the sidebar, go to **DHCP > Groups**. The page **All groups** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the group of your choice, click on **⚙**. The properties page opens.
4. In the panel **DHCP options**, click on **EDIT**. The wizard **Configure DHCP options** opens.
5. In the drop-down list **Options category**, select the option type of your choice. The wizard refreshes.
6. Edit the option(s) of your choice.
7. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Deleting a DHCP Group

At any point you can delete a DHCP group if you no longer use it or if its configuration no longer suits your needs.

Keep in mind that you cannot delete a group if it contains statics. You must first remove the statics from the group and then follow the procedure below.

To delete a DHCP group

1. In the sidebar, go to **DHCP > Groups**. The page **All groups** opens.
2. On the right-end side of the menu, click on **V4** to display the DHCPv4 groups or on **V6** to display the DHCPv6 groups.
3. Filter the list to display the group you want to delete.
4. Tick the group(s) you want to delete.
5. In the menu, click on **🗑 Delete**. The wizard **Delete** opens.
6. Click on **OK** to complete the operation. The report opens and closes. The selected group is no longer listed.

Chapter 28. Managing Dynamic Addressing

DHCP dynamic addressing relies on [ranges](#) and [leases](#).

You can also restrict access using ACLs, configure the PXE or prevent IP duplication, as detailed in the sections [Restricting Access](#), [Configuring the PXE](#) and [Preventing IP Address Duplication](#).

Managing DHCP Ranges

Ranges must be declared in SOLIDserver for dynamic addressing. A DHCP range is a contiguous suite of valid IP addresses which are available for lease to client computers on a particular scope. A range belongs to just one DHCP scope, and contains the leases of the dynamic addresses. Several ranges can be defined in the same scope if they do not overlap each other.

EfficientIP DHCP servers manage the statics with IP address like leases. Therefore the statics added also create a lease whenever the MAC address declared is active on the network, these leases can belong to your ranges and are listed on the page *All leases*. If you configured the IPAM and DNS replication, they also create DNS entries. For more details, refer to the section [Adding DHCPv4 Statics](#).

Browsing DHCP Ranges

The DHCP ranges are the third level of a DHCP dynamic addressing organization. They belong to a scope and contain the leases.

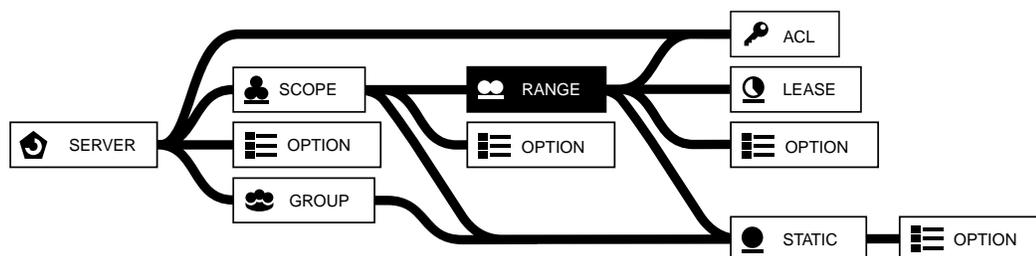


Figure 28.1. The range in the DHCP hierarchy

Browsing the Ranges Database

To display the list of DHCP ranges

1. In the sidebar, go to [DHCP > Ranges](#). The page **All ranges** opens.
2. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.
3. To display the list of ranges of a DHCP scope, in the **Scope** column, click on the name of the DHCP scope of your choice. The page refreshes.

To display a DHCP range properties page

1. In the sidebar, go to [DHCP > Ranges](#). The page **All ranges** opens.

2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the range of your choice, click on . The properties page opens.

Customizing the Display on the Page All Ranges

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that in IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the DHCP Range Statuses

The column **Status** provides information regarding ranges you manage.

Table 28.1. DHCP range statuses

Status	Description
 OK	The range is configured
 <i>Delayed create</i>	The creation or update is delayed due to a scheduled configuration of the server. The creation is automatically done after maximum of 1 minute.
 <i>Delayed delete</i>	The deletion is delayed due to a scheduled configuration of the server. The deletion is automatically done after maximum of 1 minute.

Adding a DHCP Range

The addition of a new range provides free addresses to DHCP clients. You can add an IPv4 or an IPv6 range from the page *All ranges*, each range is defined by its first and last address.

Note that you can also import ranges, for more details refer to the section [Importing Ranges](#).

Keep in mind that if you or your administrator configured IPAM to DHCP advances properties, new ranges may be added for every pool created in the IPAM. For more details, refer to the chapter in the section IPAM Advanced properties of the chapter [Managing Advanced Properties](#).

Adding a DHCPv4 Range

You can add as many DHCPv4 ranges as you need. Note that by default:

- The DHCP range addition wizard provides a page dedicated to configuring Access Control Lists (ACL) for EfficientIP DHCP servers. Keep in mind that **the order of the elements of the range's ACL list is important as each restriction or permission is reviewed following the order you set in the list**.
- You cannot add a range containing more than 1 million addresses. To edit this limit, refer to the procedure [To edit the registry key that defines the ranges maximum size](#).

To add a DHCPv4 range

1. In the sidebar, go to  **DHCP > Ranges**. The page **All ranges** opens.

2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. In the menu, click on **+ Add**. The wizard **Add a DHCP range** opens.
4. In the list **DHCP Server**, select a server.
5. Click on **[NEXT]**. The scope selection page opens.
6. In the list **DHCP Scope**, select a scope.
7. Click on **[NEXT]**. The next page opens.
8. If you or your administrator created classes at the range level, in the list **DHCP range class**, select a class or *None*. Click on **[NEXT]**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. Configure the DHCP range parameters following the fields below:

Table 28.2. DHCPv4 range parameters

Field	Description
Start address	First address of the range.
End address	Last address of the range. If you edit this address, the field <i>Size</i> is automatically updated.
Size ^a	Number of addresses in the range. If you edit this field, the field <i>End address</i> is automatically updated.
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

^aNote that you cannot add a DHCP range containing more than a million IP addresses.

10. If you are adding a range on an EfficientIP DHCP server, click on **[NEXT]**. The ACLs configuration page opens, depending on the classes configured by your administrator.

You can set the ACL configuration of your choice, using Specific and/or General ACLs.

- a. In the field **Specific ACL**, configure ACLs using the table below.

Table 28.3. Specific ACL configuration fields

Field	Description
Specific ACL	Type in the first letters of the name of an existing ACL, the auto-completion provides a list matching your search. Select the ACL of your choice.
Allow	This box allows to grant (tick) or deny (do not tick) access to the selected ACL. The permission you choose edits the content of the field <i>ACL</i> .
ACL	This field displays the configuration for the specified ACL: <i>deny members of "<ACL>"</i> or <i>allow members of "<ACL>"</i> . Once the configuration suits your needs click on + . The configuration is moved to the list <i>DHCP range ACL</i> .

- b. In the field **General ACL**, configure ACLs using the table below.

Table 28.4. General ACL configuration fields

Field	Description
General ACL	Select <i>unknown clients</i> , <i>known clients</i> , <i>all clients</i> or <i>dynamic bootp clients</i> .
Allow	This box allows to grant (tick) or deny (do not tick) access to the selected ACL. The permission you choose edits the content of the field <i>ACL</i> .
ACL	This field displays the configuration for the specified ACL: <i>deny "<ACL>"</i> or <i>allow "<ACL>"</i> . Once the configuration suits your needs click on  . The configuration is moved to the list <i>DHCP range ACL</i> .

- c. In the list **DHCP range ACL**, all the configured ACLs are listed. Set the ACL order according to your needs: select an ACL and move it up or down the list using  and .

To delete an ACL, select it and click on .

11. Click on  to complete the operation. The report opens and closes. The ACLs are listed in the ACL panel of the range properties page.

If you want to add ranges containing more than a million addresses, you must edit the dedicated registry database key.

To edit the registry key that defines the ranges maximum size

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in *module.dhcp.range_max_size*. Only this key is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in the value of your choice. The default value is *1000000* IPv4 addresses. For performance purposes, we strongly advise against setting a value greater than *7000000*.
6. Click on  to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Adding a DHCPv6 Range

You can add as many DHCPv6 ranges as you need. Note that the ACL configuration is not available on IPv6 ranges.

To add a DHCPv6 range

1. In the sidebar, go to  **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. In the menu, click on  **Add**. The wizard **Add a DHCPV6 range** opens.
4. In the list **DHCP Server**, select a server.

5. Click on **[NEXT]**. The next page opens.
6. In the list **DHCP Scope**, select a scope.
7. Click on **[NEXT]**. The next page opens.
8. If you or your administrator created classes at the range level, in the list **DHCP range class**, select a class or *None*. Click on **[NEXT]**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

9. Configure the DHCPv6 range following the table below:

Table 28.5. DHCPv6 range parameters

Field	Description
Start address	Type in the range start address, it edits the content of the field <i>Size</i> . By default, the field automatically displays the selected scope start address. This field is required.
End address	Type in the range end address, it edits the content of the field <i>Size</i> . By default, the field automatically displays the selected scope end address. This field is required.
Size	Type in the number of addresses you want in the range. The number you type in modifies the range end address.
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

10. Click on **[OK]** to complete the operation. The report opens and closes. The range is listed.

Editing a DHCP Range

Once created, DHCPv4 ranges can be modified as far as their ACL and size are concerned.

Editing a Range Properties

You can edit a DHCPv4 range advanced properties and ACLs once created from its properties page.

You cannot edit a DHCPv6 range. The properties page only display all the information available.

To edit a DHCPv4 range

1. Open the properties page of the range of your choice. For more details, refer to the procedure [To display a DHCP range properties page](#).
2. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DHCP range** opens.
3. You cannot edit the range *Start address*, *End address* and *Size*.

- Depending on the administrator configuration, you may be able to configure advanced properties according to the table below.

Table 28.6. Advanced properties configuration

Field	Description
Advanced Properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DHCP properties are detailed in the section Configuring DHCP Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The page **ACL configuration** opens¹.
- Configure or edit ACL following the table below:

Table 28.7. ACL configuration options

Field	Description
Specific ACL	Type in the name of the ACL you want to grant or deny access at this level in the DHCP. The auto-completion allows you to find them more easily.
General ACL	You can create exceptions to apply to <i>unknown clients</i> , <i>known clients</i> , <i>all clients</i> or <i>dynamic bootp clients</i> .
Allow	Tick this box if you want to allow the parameters set up in the fields <i>Specific ACL</i> and <i>General ACL</i> . If you do not tick it, what you specified in those fields is denied.
ACL	This field displays each ACL section configuration. It is gray by default because its content depends on what you configured above. Once your configuration is visible and suits your needs, click on + . The configuration is then listed the DHCP range ACL list.
DHCP range ACL	This list sums up all the ACLs configured through the wizard.

- Click on **OK** to complete the operation. The report opens and closes. The modifications are visible in the panels *Advanced properties* and/or *ACL*.

Resizing a Range

With DHCPv4, you can resize ranges. Basically, you can edit the range start and/or end address so that it includes more or less addresses. This shift in addresses is only possible if the addresses included or excluded are not already used or part of another range.

So if your range is *192.168.0.10-192.168.0.125* you can decide to resize to *192.168.0.100-192.168.0.105* indicating a start address shift of "90" and an end address shift of "-20".

To resize a DHCPv4 range

- In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
- On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.

¹For more details regarding the ACLs, refer to the section [Restricting Access](#).

3. Filter the list if need be. For more details, refer to the procedure [To display the list of DHCP ranges](#).
4. Tick the range(s) you want to resize.
5. In the menu, select  **Edit > Resize DHCP ranges**. The wizard **Resize ranges** opens.
6. In the field **Start address shift**, type in the positive or negative shift for the range start address that suits your needs. If you type in 0 (zero), the address stays the same.
7. In the field **End address shift**, type in the positive or negative shift for the range end address that suits your needs. If you type in 0 (zero), the address stays the same.
8. Click on to complete the operation. The report opens and closes. The new range(s) size is visible.

Replicating Range Data in the IPAM

At range level, the option *IPAM replication* allows to replicate range data in IPAM pools. Before replicating your DHCP data in the IPAM, keep in mind that:

- The range data is replicated in pools of the IPAM.
- The option *IPAM replication* behaves as follows:
 - If the range belongs to a scope for which the IPAM replication has been set, its data is replicated as a pool that belongs to the same terminal network.
 - If the range belongs to a scope for which no replication has been set, its data is replicated as a pool in the first terminal network that can receive it.
 - You can replicate the range data in the IPAM if:
 - No pool exists yet: a read-only pool is created and named *DHCP*.
 - An existing matching the range size already exists. If the pool was not in read-only it is marked read-only, the pool is renamed *DHCP*.
 - You cannot replicate a range data in the IPAM if an existing pool already manages some of the IP addresses of the range you selected.
- The option *IPAM replication* is independent at range level, replicating scopes does not automatically replicates to the ranges they contain. You must select ranges and execute the option to update the IPAM pools with range data.

For more details regarding scope replication, refer to the section [Defining a Specific IPAM Space for a DHCPv4 Scope](#).

To replicate range data in the IPAM

1. In the sidebar, go to  **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. Tick the range(s) of your choice.
4. In the menu, select  **Edit > IPAM replication**. The wizard **IPAM replication** opens.
5. Click on to complete the operation. The report opens and closes.

The matching pool is visible on the IPAM page **All pools**.

Configuring DHCP Options at Range Level

You can configure DHCP options at range level, individually or in bulk. Keep in mind that:

- You can only edit DHCP options at range level on DHCPv4 servers.
- The DHCP options of a range may be inherited from its server or scope.
- The options set at range level override the options set on its container.
- All the DHCP options you configure at range level are inherited by all leases it delivers.
- The DHCP options configuration at range level is not supported on Microsoft DHCP servers.

For more details regarding the DHCP options configuration, refer to the chapter [Configuring DHCP Options](#) and/or to the appendix [DHCP Options](#).

Editing the DHCP Options of a Range

From the properties page of a range, you can set DHCP options.

To edit DHCP range options

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. At the end of the line of the range of your choice, click on **[⚙️]**. The properties page opens.
4. In the upper right corner, click on **[☰]**.
5. In the panel **DHCP Options**, click on **[EDIT]**. The wizard **Configure DHCP options** opens.
6. In the drop-down list **Options category**, select the option type of your choice. The wizard refreshes.
7. Edit the option(s) of your choice.
8. Click on **[OK]** to complete the operation. The report opens and closes. The changes are listed in the panel.

Performing Option Changes on Several Ranges At Once

From the page All ranges, you can set, replace or delete DHCP options on all the ranges you select at once.

To add a DHCP option to one or several ranges

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the range(s) of your choice.
5. In the menu, select **☑️ Edit > Option > Add**. The wizard **Add DHCP range options** opens.
6. In the drop-down list **Option name**, select an option.
7. In the field **Value**, type in the relevant value.

8. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the range properties page, the panel **DHCP options** lists the new DHCP option and its value.

To edit the value of a DHCP option on one or several ranges

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the range(s) of your choice.
5. In the menu, select **Edit > Option > Replace**. The wizard **Replace DHCP range options** opens.
6. In the drop-down list **Option name**, select the option which value you want to replace.
7. In the field **Replace**, specify the value you want to change.
8. In the field **By**, type in the new option value.
9. Click on **OK** to complete the operation. The report opens and closes, the page refreshes. On the range properties page, the panel **DHCP options** displays the new value of the DHCP option.

To remove a DHCP option from one or several ranges

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **V4**. The page refreshes and the button turns black.
3. Filter the list if need be.
4. Tick the range(s) of your choice.
5. In the menu, select **Edit > Option > Delete**. The wizard **Delete DHCP range options** opens.
6. In the drop-down list **Option name**, select an option.
7. In the field **Option value filter**, type in the option value.
8. Click on **OK** to complete the operation. The report opens and closes. On the range properties page, the panel **DHCP options** no longer displays the DHCP option.

Deleting a DHCP Range

When a network is no longer used, or whenever you wish, you can delete an existing range. The deletion procedure is identical for IPv4 and IPv6 ranges.

Before deleting an existing range, remember to create a new one using a different range of addresses.

To delete a DHCP range

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.

3. Filter the list if need be. For more details, refer to the procedure [To display the list of DHCP ranges](#).
4. Tick the range you want to delete.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on to complete the operation. The report opens and closes. The range is no longer listed.

Managing DHCP Leases

When it comes to dynamic addressing, the lease is the lowest level of the DHCP hierarchy. A lease corresponds to one IP address, listed in the IPAM module. Just like its name indicates it, a lease is limited in time. When a client requests an IP address to a DHCP server, the server delivers an IP address that is part of the scope that listens to the network area where the client asked for an address. Which is why it is important to properly set up the DHCP server. Once you created at least one scope and one range in a DHCP server you are able to deliver leases.

Keep in mind that **if the appliance time is incorrect, you cannot retrieve any leases**. For more details, refer to the section [Configuring NTP Servers](#).

With SOLIDserver in DHCPv4, the maximum lease time is 24 hours (86400 seconds). By default, the lease time is of 12 hours (43200 seconds). You can obviously change these parameters either one a particular lease individually or at the range, scope or server level. As for DHCPv6, you can configure the leases only at the server or scope level.

EfficientIP DHCP servers manage the statics with IP address like leases. Therefore in IPv4:

- On the DHCP page *All leases*, all the clients are listed whether they requested access to the server on parts of the network - defined through your scopes and ranges - or whether they were identified through their MAC address. For more details, refer to the section [Adding DHCPv4 Statics](#).
- On the IPAM page *All addresses*, if the IP address leased is also managed in the IPAM, the columns *Type* and *Status* of the IP address reflect its use in the DHCP. For more details, refer to the section [Understanding the IP Address Type and Status](#).

Note that **once your leases have expired, are released or are deleted**, they are no longer listed on the page *All leases*: **they are moved to the page Lease history**. For more details, refer to the section [Tracking Leases](#).

Browsing the Leases

The leases are the last level of the DHCP dynamic addressing hierarchy.

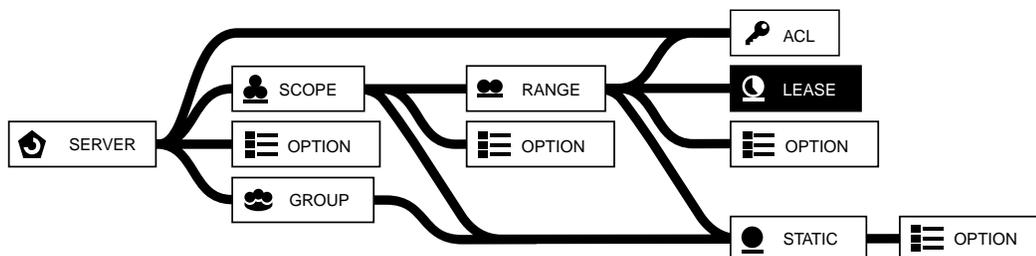


Figure 28.2. The lease in the DHCP hierarchy

Browsing the Leases Database

To display the list of DHCP leases

1. In the sidebar, go to **DHCP > Leases**. The page **All leases** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To display the list of leases of a DHCP range, in the column **Range**, click on the name of the DHCP range of your choice. The page refreshes.

To display a DHCP lease properties page

1. In the sidebar, go to **DHCP > Leases**. The page **All leases** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Filter the list if need be.
4. At the end of the line of the lease of your choice, click on **ⓘ**. The properties page opens.

Customizing the Display on the Page All Leases

Users of the group *admin* can create customized column layouts. The button **☰ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Note that the columns on this page allow to display all the current IPv4 or IPv6 lease information, like their **Status**, **IP address**, **MAC address**, **MAC vendor**, **Start** and **End** time and date; or even **OS name**, **MAC type**, **Circuit ID**, **Client DUID**, **Remote ID**, **Multi-status**, **Vendor ID**, etc. To display the lease logs, refer to the section [Tracking Leases](#).

Keep in mind that in IPv6 you can display colored labels above parts of the IP addresses listed. It allows to differentiate at a glance your containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Understanding the DHCP Lease Statuses

The column **Status** provides information regarding the leases you manage.

Table 28.8. DHCP lease statuses

Status	Description
 <i>OK</i>	The lease is configured
 <i>Delayed create</i>	The creation or update is delayed due to a scheduled configuration of the server. The creation is automatically done after maximum of 1 minute.
 <i>Delayed delete</i>	The deletion is delayed due to a scheduled configuration of the server. The deletion is automatically done after maximum of 1 minute.

Defining the Lease Duration

When a DHCP client requests an IPv4 or an IPv6 address, it may suggest a lease duration in the DHCPDISCOVER message. If the client requests a particular lease duration, the server makes sure the requested lease time is within a range specified by the *min-lease-time* and *max-lease-time* parameters. If the requested lease time is not within the specified range, it is set to

the value of *min-lease-time* if it is too short or to the value of *max-lease-time* if it is too long. If the client does not request a specific lease duration, the lease duration specified in the *default-lease-time* is used, and the same limits are applied.

EfficientIP DHCP server allows administrators to specify a default lease duration, a minimum lease duration, and a maximum lease duration as defined below:

- **default-lease-time** specifies the duration of the lease that the DHCP server assigns if the client requesting the lease does not ask for a specific expiration time.
- **minimum-lease-time** duration is used to force the DHCP client to take a longer lease than the lease duration that it requests.
- **maximum lease-time** duration is used to define the longest lease that the DHCP server can allocate. If a DHCP client asks for a longer lease than the maximum lease duration, then the server limits the lease to the maximum lease duration.

Note that the **maximum lease times does not apply to dynamic BOOTP leases**. These leases are not specified by the client and can exceed the maximum lease time configured.

You can set up the lease duration at server, scope, range, group and static level in IPv4 and at server and scope level in IPv6. You can also configure a DHCP class to set the lease duration.

DHCP lease duration is a topic of discussion among network administrators. Some use a lease time of 6 months, some use lease time of 5 minutes. The right lease duration depends on each network's context. Default lease duration on EfficientIP DHCP server is 12 hours. You can change this default according to your requirements and set leases time at different levels, based on different factors. You can set a default lease time at the server, scope, range, group, DHCP class, or static level of the EfficientIP DHCP organization.

To configure lease duration in DHCPv4

1. In the sidebar, go to **DHCP > Servers, Groups, Scopes, Ranges** or **Statics** depending on your needs. The page opens.
2. Filter the list if need be.
3. At the end of the line of the object of your choice, click on **⌵**. The properties page opens.
4. In the panel **DHCP options**, click on **[EDIT]**. The wizard **Configure DHCP options** opens.
5. In the field **Default lease time**, you can set a default lease time in seconds. The lease time is respected unless the client specifies another one when requesting a lease.
6. In the field **Max lease time**, you can set the maximum lease time in seconds.
7. In the field **Min lease time**, you can set the minimum lease time in seconds.
8. Click on **[OK]** to complete the operation. The report opens and closes. The edited information is now listed in the panel.

To configure lease duration in DHCPv6

1. In the sidebar, go to **DHCP > Servers, Scopes** or **Statics** depending on your needs. The page opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. Filter the list if need be.

4. At the end of the line of the object of your choice, click on . The properties page opens.
5. In the panel **DHCP options**, click on . The wizard **Configure DHCP options** opens.
6. In the field **Default lease time**, you can set a default lease time in seconds. The lease time is respected unless the client specifies another one when requesting a lease.
7. In the field **Max lease time**, you can set the maximum lease time in seconds.
8. In the field **Min lease time**, you can set the minimum lease time in seconds.
9. Click on  to complete the operation. The report opens and closes. The edited information is now listed in the panel.

Releasing Leases

In case of ranges overloading, the lease release feature can be helpful in order to punctually free a critical case. This operation asks the DHCP server to simulate a DHCP release.

Releasing leases should not be done on a daily basis to resolve a lack of free space in a range, in this case it is best to extend the range capacity as soon as possible.

Besides, keep in mind that **a lease deletion can create IP addresses overlapping**. Before proceeding with the lease deletion, make sure that the impacted DHCP client is not liable to connect to the network where the addresses were deleted.

To delete a DHCP lease

1. In the sidebar, go to  **DHCP > Leases**. The page **All leases** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Filter the list to find the lease(s) you want to delete.
4. Tick the lease(s) to be deleted.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on  to complete the operation. The report opens and closes. The selected leases are no longer listed.

Converting IPv4 Leases into Statics

SOLIDserver allows you to convert leases to static reservations in order to register a MAC address in one or several DHCP servers. No matter how you distributed the leases among servers, the moment you convert a lease into a static, all the servers on the failover are notified and can grant the MAC address the DHCP options you configured.

It is not possible to convert an IPv6 lease into a static but you can create IPv6 statics.

When you convert a lease to static, a static reservation is created with the same name as the lease. This reservation can have an IP address or not:

- **Converting into a static without IP address:** the MAC address of the lease now connects to the first available IP address on the network - no matter the server, scope or range managing it. The purpose of the conversion is to configure DHCP options for the static reservation that applies to the MAC address whenever it connects to the network.
- **Converting into a static with IP address:** the MAC address of the lease always connects to the same IP address. The purpose of this conversion is to configure the same specific DHCP

options for a specific MAC address, or client, whenever they connect to the part of the network that manages their IP address.

To convert an IPv4 lease into a static

1. In the sidebar, go to **DHCP > Leases**. The page **All leases** opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Filter the list to find the lease(s) you want to convert.
4. Tick the lease(s) to be converted.
5. To convert a lease into a static without IP address, in the menu select **[edit icon] > Edit > Convert to static > Without IP address**. The wizard **Convert lease to static without IP address** opens.
6. To convert a lease into a static with IP address, in the menu, select **[edit icon] > Edit > Convert to static > With IP address**. The wizard **Convert DHCP lease to DHCP static** opens.
7. Click on **[OK]** to complete the operation. The report opens and closes.

The converted IP addresses are no longer listed on the page *All leases*. They are now on the page *All statics*, accessible via the breadcrumb.

Blacklisting Leases

Once delivered, you can blacklist a lease at any time. This converts the lease into a static without IP. From that point on, the client MAC address cannot access to the DHCP servers or the failover channel and therefore can no longer be delivered a lease.

Once a lease is blacklisted, the corresponding static without IP is immediately created. The client MAC address is saved in the DHCP server configuration as *blacklist-<MAC_address>* to ensure that any lease request is denied. This static is automatically configured with a set of ACL restrictions that prevent the connection to the server and its failover. In the meantime, the lease remains valid until it expires, the next client request for renewal is denied. Once the lease duration is up, the client MAC address is disconnected and unable to connect again.

EfficientIP DHCP servers manage the statics with IP address like leases². The static reservations create leases, identified via their MAC address, that you can also blacklist: a static without IP address automatically replaces the static with IP address you blacklisted.

To blacklist a lease

1. In the sidebar, go to **DHCP > Leases**. The page **All leases** opens.
2. Tick the lease(s) you want to blacklist.
3. In the menu, select **[edit icon] > Edit > Blacklist lease**. The report opens and closes. The lease is still visible on the page **All leases** and disappears once it has expired. On the page **All statics**, every blacklisted MAC address as the following **Name**: *blacklist*.

Tracking Leases

SOLIDserver keeps track of the leases delivered by all the DHCP servers you manage. You can purge the lease history via two rules that allow to remove expired leases from the database after a certain number of days. For more details refer to the sections and .

²For more details, refer to the section [Adding DHCPv4 Statics](#).

The lease logs are available on the page *Lease history*.

To track leases

1. In the sidebar, go to **DHCP > Leases**. The page **All leases** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. On the right-end side of the menu, click on **DHCP lease history**. The page refreshes.

The page provides information on each lease through the columns:

Table 28.9. The columns available on the pages Lease tracking and Lease tracking (v6)

Field	Description
IP address	The IP address allocated.
MAC address	The MAC address of the client that used the lease.
Start	The lease allocation start time and date.
End	The lease expiration time and date.
Server	The name of the server that delivered the lease.
Client identifier ^a	The value of the option <i>client-identifier</i> sent by the client that requested the DHCPv4 lease.
Remote ID ^a	The remote ID provided by the relay agent that received the DHCPv4 lease request and sent it to the server.
Circuit ID ^a	The circuit ID provided by the relay agent that received the DHCPv4 lease request and sent it to the server.
Period ^a	The total lifespan of the DHCPv4 lease.
Name	The name of the lease.
OS name ^a	The DHCP client OS name and version of the client to which the DHCPv4 lease was allocated.
Status	The status of the lease.

^aThis column is not available for DHCPv6 leases.

Purging Expired DHCP Leases

The IPv4 lease logs are automatically erased 60 days after the leases have expired, as set in rule 012. You can change this rule configuration following the procedure below.

To edit the rule 012 that controls the automatic purge of IPv4 lease logs

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #** search field, type in *012*. The rule *Purge DHCP leases history* is listed.
4. In the column **Instance**, click on *auto_purge_histo_dhcplease*. The rule properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a rule** opens.
6. Click on **NEXT**. The page **Rule filters** appears.
7. If you want to schedule the purge, configure the fields according to the table below:

Table 28.10. Rule filters parameters

Field	Description
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, 23 is selected.
Minute	Select a period of time, minutes-wise. By default, 15 is selected.

8. Click on **[NEXT]**. The **Rule parameters** appears.
9. In the field **Number of days**, type in the number of days beyond which an expired lease should be removed from the logs. By default, the value is 60.
10. Click on **[OK]** to complete the operation. The report open and closes. The rule properties page is visible again.

You can at any time disable this rule.

To disable the rule 012

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #** search field, type in 012.
4. Tick the rule 012 and, in the menu, select **✎ Edit > Disable**. The menu **Disable** opens.
5. Click on **[OK]** to complete the operation. The wizard closes, the page refreshes. The rule is **⊗ Disabled**.

Purging Expired DHCPv6 Leases

The IPv6 lease logs are automatically erased 60 days after the leases have expired, as set in rule 384. You can change this rule configuration following the procedure below.

To edit the rule 384 that controls the automatic purge of IPv6 lease logs

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #** search field, type in 384. The rule *Purge DHCPv6 leases history* is listed.
4. In the column **Instance**, click on *auto_purge_histo_dhcplease6*. The rule properties page opens.
5. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a rule** opens.
6. Click on **[NEXT]**. The page **Rule filters** appears.
7. If you want to schedule the purge, configure the fields according to the table below:

Table 28.11. Rule filters parameters

Field	Description
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.

Field	Description
Hour	Select a specific time or one of the available schedules. By default, 23 is selected.
Minute	Select a period of time, minutes-wise. By default, 15 is selected.

8. Click on **NEXT**. The page **Rule parameters** appears.
9. In the field **Number of days**, type in the number of days beyond which an expired lease should be removed from the logs. By default, the value is 60 days.
10. Click on **OK** to complete the operation. The report open and closes. The rule properties page is visible again.

You can at any time disable this rule.

To disable the rule 384

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #** search field, type in 384.
4. Tick the rule 384 and, in the menu, select **Edit > Disable**. The menu **Disable** opens.
5. Click on **OK** to complete the operation. The wizard closes, the page refreshes. The rule is **Disabled**.

Displaying the Relay Agent Information (Option 82)

To put it simply, DHCPv4 Option 82 is the *DHCP Relay Agent Information* option. The DHCP relay agent and Option 82 are defined in RFC 3046. Option 82 was designed to allow a DHCP relay agent to insert circuit specific information into a request that is being forwarded to a DHCP server. Specifically the option works by setting three sub-options: Circuit ID, Remote ID and GIADDR.

- The Circuit ID field generally contains information describing the port location that the DHCP request is coming in from. It may contain additional information that helps describe which IP address should be assigned out, such as the VLAN ID, a wireless modem or an ATM virtual circuit. This value must be unique for a particular switch or router that is providing the Relay Agent function. The value must also stay the same if modules are installed or removed in the Switch or Router that implements the Relay Agent. Therefore, having subfields representing the Module, Slot and Port is highly recommended.
- The Remote ID field is intended to carry information describing the device at the remote end of the link. However, in Ethernet systems, this is typically the MAC address of the Relay Agent. This is not particularly useful since the MAC address would change if the Relay Agent was ever replaced. Building a DHCP server database using the MAC address of the Relay Agent would require that the table be rebuilt every time one of the Relay Agents was replaced. Some vendors have modified this field to use the IP address of the Relay Agent or some other string describing the Relay Agent. This field must be unique to the entire network.
- The GIADDR (or Gateway Address) field is part of the normal DHCP message. It contains the IP address of the Relay Agent. Since IP addresses must be unique, this field is unique for the entire network.

By combining the GiAddr and the Circuit ID, a network wide unique string can be created. This string can be used for table lookup in the DHCP server. We called this string a pseudo MAC address, since most DHCP servers do a MAC to IP mapping in their databases.

In its default configuration, the DHCP Relay Agent Information Option passes along port and agent information to SOLIDserver DHCP server. It is useful in statistical analysis for instance, as it indicates where an assigned IP address physically connects to the network. It may also be used to make DHCP decisions based on where the request is coming from or even which user is making the request. For more details regarding its implementation, refer to the chapter [Configuring DHCP Options](#).

This information is only available on EfficientIP DHCP servers and is not available on the other vendors' DHCP servers. For EfficientIP servers, you can display the **Circuit ID** (DHCP lease circuit ID) and **Remote ID** (DHCP lease remote ID) columns on the page All leases. For more details, refer to the section [Customizing the List Layout](#).

The Relay Agent Information with DHCPv6

With DHCPv6, the client ID, circuit ID and remote ID are not supported. It is impossible therefore to retrieve these pieces of information separately, much less display them in a listing template on the leases page. This information might be delivered by the agent in DHCPv6 but the appliance does not retrieve it at the server level.

The equivalent of the option 82 relay agent would be the DHCPv6 option 9 (relay message option) and the option 47 (relay data option).

Restricting Access

When a DHCP client requests an IP address, SOLIDserver offers an address from a range associated with the network segment for that client. In addition to identifying DHCP clients and allocating addresses to them, you might want to identify clients for other reasons.

For instance, you can control the access to leases or restrict IP address allocation to DHCP clients the network administrators do not know. Some network areas might want to group clients in some way or you might have to allocate dynamic IP addresses for known clients on a particular network segment and for unknown clients on the same network segment but on a different IP network.

No matter the way you want to control access, you can use access control lists, or ACL.

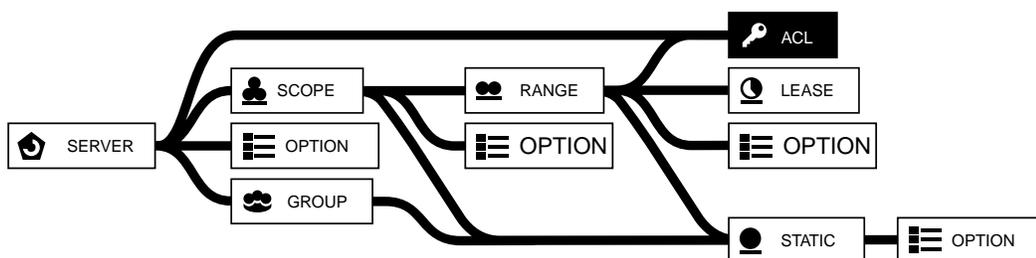


Figure 28.3. The ACL in the DHCP hierarchy

This configuration is not available on Microsoft Windows servers. Access control is only available on EfficientIP SOLIDserver appliances and on ISC DHCP delivered in EfficientIP's packages for Linux, Solaris and FreeBSD.

Granting Access to Known Clients

If you want to set up a SOLIDserver that provides dynamic IP addresses only to known clients, you first need to declare static reservations for these clients with a client identifier or a MAC ad-

dress, without specifying an IP address for them. Then, you must configure the DHCP server not to provide IP addresses to unknown clients, in order to limit access to DHCP clients for which static reservations exist. To apply this mechanism you have to setup the *Allow Known Client* ACL on the DHCP ranges.

DHCPv6 does not support ACLs configuration, you can only use access control lists on DHCPv4 objects.

To grant DHCP access only to known clients

1. In the sidebar, go to **DHCP > Ranges**. The page **All ranges** opens.
2. At the end of the line of the range of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **ACLs**, click on **EDIT**. The wizard **Edit a DHCP range** opens.
4. Click on **NEXT** to skip the range main information. The next page of the wizard appears.
5. In the section **Specific ACL**, do not modify anything.
6. In the section **General ACL**, in the drop-down list select *known clients*.
7. In the section **Allow**, tick the box. The field **ACL** displays *allow known clients*.
8. Click on **ⓘ**. *allow known clients* is now preset in the list **DHCP range ACL**.
9. Click on **OK** to complete the operation. The report opens and closes. The modification is visible in the panel **ACL**.

Restricting Access Using ACLs

SOLIDserver offers a construct called ACL that you can use to group DHCPv4 clients in a more general manner than you can do with a static reservation. Like static reservations, ACLs can be used as a client membership to control how addresses are allocated. DHCP clients become members of ACLs either because they match an ACL matching rule or because they match an entry of that ACL. ACLs can be applied to allow or deny the dynamic allocation from a range of IP addresses. Once an ACL is defined it can be used several times to restrict the access to a range.

The DHCP module provides two different lists regarding ACLs:

- The page **All ACLs** is accessible in the breadcrumb additional pages of the pages All servers, scopes, ranges and leases.
- The page **ACL Entries** that is only accessible through the breadcrumb on the page All ACLs.

Adding, Editing and Copying ACLs

From the page All ACLs you can add ACLs that grant or deny access to the DHCPv4 servers or smart architectures of your choice. The ACL is a succession of checks that ultimately make sure that all the parameters you want or refuse from your DHCPv4 clients toward the DHCP server or smart architecture of your choice are respected. There are a number of predefined ACLs available upon creation if you want to apply specific behaviors or simply reuse the syntax and configure a custom-made ACL. Among them, only the MAC address checks a list of data rather than parameters.

To add an ACL

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.

2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. In the menu, click on **+ Add**. The wizard opens.
5. In the list **DHCP server**, select one of your DHCPv4 servers.
6. Click on **NEXT**. The page **DHCP ACL parameters** opens.
7. In field **ACL name**, type in the name of the ACL to be created.
8. In the drop-down list **Predefined ACL**, you can select one of the available ACLs. The ACL syntax is displayed in the field **ACL rule** and can be edited. By default, *None* is selected and nothing is displayed in the field *ACL rule*.
9. In the field **ACL rule**, type in or modify the syntax if need be.
10. Click on **OK** to complete the operation. The report opens and closes. The ACL is listed.

Once added, the ACL can be configured to be even more efficient. For instance, if you used the MAC address ACL or an ACL comparing a list of information, you can define an ACL Entry to set up the corresponding parameters, and make sure, for example, that the access list is granted or denied only to the MAC address of your choice.

To edit an ACL

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. Filter the list if need be.
5. At the end of the line of the ACL of your choice, click on **ⓘ**. The properties page opens.
6. In the panel **Main properties**, click on **EDIT**. The wizard **DHCP ACL parameters** opens.
7. Edit the fields **Predefined ACL** and **ACL rule** according to your needs. Note that you cannot rename an ACL.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

At any time, you can copy an ACL from one server to the other. This ACL duplication copies the ACL entries as well. However, once copied, you still have to assign each new ACL in the target server to use it.

Keep in mind that if your physical server is managed via a smart, only the ACL created on the smart can be duplicated.

To copy an ACL

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. Filter the list if need be.
5. Tick the ACL(s) you want to copy.
6. In the menu, select **↗ Edit > Migrate**. The wizard **Copying ACLs** opens.

7. In the drop-down list **Target server**, select the server or smart architecture of your choice.
8. Click on **OK** to complete the operation. The report opens and closes. On the unfiltered **All ACLs** list, the duplicate ACLs are listed.

If you migrate an ACL to a smart architecture that manages physical servers, the ACL is copied to the smart and then pushed to the physical server: it stays in *Delayed create* until it is successfully pushed.

Adding ACL Entries

Once you added an ACL, you can add ACL entries to the ACL to define the rule that governs the ACL you are adding. Note that you can only add or delete ACL Entries, you cannot edit them even from their properties page.

To add an ACL Entry

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. In the breadcrumb, click on **ACL Entries**. The page **ACL Entries** opens.
5. In the menu, click on **+ Add**. The wizard **Add an ACL entry** opens.
6. In the list **DHCP server**, select one of your DHCPv4 servers.
7. Click on **NEXT**. The next page of the wizard opens.
8. In the list **DHCP ACL**, select the ACL of your choice.
9. Click on **NEXT**. The last page of the wizard opens.
10. In the field **ACL Entry**, type in your condition following the a MAC format of two hexadecimal characters separated by a semi-colon. If your ACL is a matching MAC address, type in the matching MAC address.
11. Click on **OK** to complete the operation. The report opens and closes. The ACL Entry is listed, it named after the server it belongs to and its value matches what you typed into the ACL Entry field.

Configuring the PXE

The PXE (Preboot eXecution Environment) is used to boot hosts using a network interface independently of available data storage devices or installed operating systems. The PXE protocol is a combination of DHCPv4 and TFTP protocol. Note that there is no Preboot eXecutable Environment boot standard for IPv6 yet.

DHCP is used to locate the appropriate boot server or servers, with TFTP used to download the initial bootstrap file. After it downloads the file, the host reboots and sends another IP address request. When such a PXE client starts up, it first requests an IP address in order to download the file it needs to boot.

The client, wishing to remotely boot an operating system image, broadcasts a DHCPDISCOVER packet as per the DHCP protocol. This packet is transmitted to acquire an IP address. The client also sends PXE protocol specific DHCP option 60 (Vendor Class Identifier) along with this packet. The DHCP server responds to the above DHCPDISCOVER packet by sending a DH-

CPOFFER packet that contains the IP Address allocated to the client. In a PXE remote boot, the DHCP server also sends:

- A **special tag** (option 60, with the value set to the string "PXEClient") to identify that it is capable of configuring a PXE client.
- The **next server** to specify the server host address from which the initial boot file is to be loaded.
- The **filename** to specify the name of the initial boot file to be loaded by a DHCP client.

The client downloads the executable file using either standard TFTP (port69) or MTFTP (port assigned in Boot Server Ack packet). The file downloaded and the placement of the downloaded code in memory is dependent on the client's CPU architecture. After it downloads the boot file, the client reboots and sends a new DHCPDISCOVER.

You can set a different lease time for PXE boot requests to manage your dynamic ranges better. The DHCP server can allocate an IP address with a shorter lease time to hosts that send PXE boot requests in order to release IP addresses faster.

Necessary Parameters for PXE

Usually, to implement the PXE protocol, DHCP options and/or BOOTP parameters must be configured:

- **Next-server** (BOOTP parameter) specifies the host address of the server from which the initial boot file (specified in the filename statement) is to be loaded. The value of this option should be a numeric IP address. If no next-server parameter applies to a given client, the DHCP server IP address is used.
- **TFTP-server-name** (DHCP option #66) is used to identify a TFTP server when the Next-server (BOOTP parameter) field in the DHCP header has been used for DHCP options.
- **Filename** (BOOTP parameter) specifies the name of the initial boot file to be loaded. The value of this option should be the name of a file that is recognizable to whatever file transfer protocol the client is expected to use to load the file. Some clients might prefer to receive this information in the *bootfile-name* option.
- **Bootfile** (DHCP option #67) specifies the name of the boot file to be used when the file field is used to carry options.

These options can be configured at multiple levels: server, scope, static reservation, DHCP group or dynamic range.

The PXE parameters configuration **only applies to DHCPv4**. For now, it is impossible to set them with IPv6 addressing.

To configure the next-server and the filename options in DHCPv4

1. In the sidebar, go to **DHCP > Servers, Groups, Scopes, Ranges** or **Statics** depending on your needs. The page opens.
2. At the end of the line of the object of your choice, click on **⚙**. The properties page opens.
3. In the panel **DHCP options**, click on **[EDIT]**. The wizard **Configure DHCP options** opens.
4. In the drop-down list **Option category**, select the **BootP Compatible** option. The two options: *next-server* and *filename* are listed among the options.

5. In the field **next-server**, type in the IP address of the server from which the initial boot file should be loaded.
6. In the field **filename**, type in the name of the initial boot file to be loaded.
7. Click on to complete the operation. The report opens and closes. The modifications are listed in the panel *DHCP options*.

Duplicated Lease with PXE

The PXE client uses two stages in its IP address request. The first is done by the hardware firmware, and the second one by the operating system. On some configuration the hardware can request IP address by using DHCP parameters that differ from the operating system, and then have two different DHCP leases for the same device. For instance, the first DHCP lease is delivered by the PXE stage by using the MAC address as lease identifier, and the operating system receives another DHCP lease based on a client identifier (sent by the client) instead of the MAC address. In this case the DHCP server believes it negotiates IP addresses for two different clients, one based on its MAC address and the other one on its client identifier.

To avoid this issue, SOLIDserver manages leases by setting a different lease time for PXE boot request. SOLIDserver allows you to allocate an IP address with a shorter lease time to hosts that send PXE boot requests, so IP addresses are not leased longer than necessary. By default the lease duration for PXE client is set to 5 minutes (300 seconds). It can be changed by following the next procedure.

To change the lease time for PXE client

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. In the search engine of the column **Name**, type in *PXE*. The PXE clients ACLs are listed.
5. At the end of the line of the ACL of your choice, click on . The properties page opens.
6. In the panel **DHCP options**, click on . The wizard **Configure DHCP options** opens.
7. In the drop-down list **Option category**, select *Most used options*.
8. In the fields **Default lease time** and **Max lease time**, type in the durations of your choice. These values are in seconds, by default they are set to 300 seconds (5 minutes)
9. Click on to complete the operation. The report opens and closes. The modifications are listed in the panel *DHCP options*.

Preventing IP Address Duplication

The *ping check* feature tells the DHCP server whether to send a ping request to check an IP address before offering it to a DHCP client using either IPv4 or IPv6. The ping check feature can protect the DHCP against address overlapping.

When the DHCP server is considering dynamically allocating an IP address to a client, it first sends an ICMP echo request (a ping) to the address being assigned. It waits for a second, and if no ICMP echo response has been heard, it assigns the address. If a response is heard, the lease is abandoned, and the server selects another free IP address and sends it a ping. The DHCP server continues this process until it finds an IP address that does not respond to the ping.

The DHCP server then sends a *DHCPOFFER* message with the unused IP address to the DHCP client.

To enable the ping check with DHCPv4

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the DHCPv4 server or smart architecture of your choice, click on **⚙**. The properties pages opens.
3. In the panel **DHCP options**, click on **[EDIT]**. The wizard **Configure DHCP options** opens.
4. In the drop-down list **Option category**, select *Server parameters*.
5. In the drop-down list **Ping check**, select *Yes*.
6. In the field **Ping timeout**, you can set up a timeout if necessary.

If the DHCP server determines that it should send an ICMP echo request (a ping) because the ping-check statement is true, ping-timeout allows you to configure how many seconds the DHCP server should wait for an ICMP Echo response to be heard, if no ICMP Echo response has been received before the timeout expires, it assigns the address. If a response is heard, the lease is abandoned, and the server does not respond to the client. If no value is set, the ping-timeout is of 1 second by default.

7. Click on **[OK]** to complete the operation. The report opens and closes. The modifications are listed in the panel *DHCP options*.

With DHCPv6, the procedure is similar. Only a few wizard-related steps change.

To enable the ping check with DHCPv6

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the DHCPv6 server or smart architecture of your choice, click on **⚙**. The properties pages opens.
3. In the panel **DHCP options**, click on **[EDIT]**. The wizard **Configure DHCP options** opens.
4. In the drop-down list **Ping check**, select *Yes*.
5. In the field **Ping timeout**, you can set up a timeout if necessary.

If the DHCP server determines that it should send an ICMP echo request (a ping) because the ping-check statement is true, ping-timeout allows you to configure how many seconds the DHCP server should wait for an ICMP Echo response to be heard, if no ICMP Echo response has been received before the timeout expires, it assigns the address. If a response is heard, the lease is abandoned, and the server does not respond to the client. If no value is set, the ping-timeout is of 1 second by default.

6. Click on **[OK]** to complete the operation. The report opens and closes. The modifications are listed in the panel *DHCP options*.

Chapter 29. Managing Failover Channels

SOLIDserver allows you to display all the failover channels for the DHCP smart architectures it manages. The page All failover channels provides you with detailed information on all the failover channels of the smart architectures you manage.

Note that if you manage servers from smart architectures you can associate a failover channel with one or several scopes at once from the page All scopes. For more details, refer to the section [Defining a Specific Failover Channel for a Scope](#).

DHCP Failover Principles and Operational States

The DHCPv4 synchronization mechanism is called failover because it was initially intended to provide a way for one DHCP server to act as a primary server and for a second DHCP server to act as a backup. In most of the basic failover configurations, the secondary server does not reply to the DHCP client requests when it is in contact with the primary, it simply synchronizes updates from the primary. In an EfficientIP DHCP configuration, both the primary and secondary servers provide simultaneously the DHCP service by default. You can change this configuration and make sure that only the primary or the secondary server responds. For more details regarding the Master/backup and Load balancing configurations of the failover, refer to the section [Operating in Normal State](#).

Keep in mind that the **failover mechanism is not available when it comes to IPv6 addressing**.

DHCP Safe Failover Principles

The failover based on load balancing involves three principles:

1. The primary and the secondary failover servers divide the dynamic ranges of free addresses that they have to serve into free and backup addresses. Free addresses are available for the primary server to allocate to its clients and backup addresses are available for the secondary server to allocate to its clients.
2. Until the servers have not exchanged leases allocation details, they can still allocate or renew leases within the range of addresses they manage but the lease time always corresponds to the Maximum Client Lead Time (MCLT). By default, it is set to 1 hour. Therefore, as long as the servers do not communicate or have not exchanged information - during the first allocation, in communications-interrupted... - all leases are set to 1 hour.
3. In normal operation, an address that has been assigned to one client cannot be assigned to another client unless both DHCP servers agree that the first client is no longer using it.

By default, the failover is based on load balancing, it is *Balanced*, on One-to-One and One-to-Many smart architectures.

DHCP Failover Operational States

There are several DHCP operational states in the failover protocol: [Operating in Normal State](#) , [Operating in Communications-interrupted State](#) and [Operating in Partner-down State](#) .

In the GUI, the page All failover channels provides the column *State* that includes detailed information regarding each state. For more details regarding this column, refer to the table [The different failover states](#)

Operating in Normal State

In *Normal* state, only one server responds to the messages sent by DHCP clients. The failover configuration you chose when setting up the failover defines which server responds to clients.

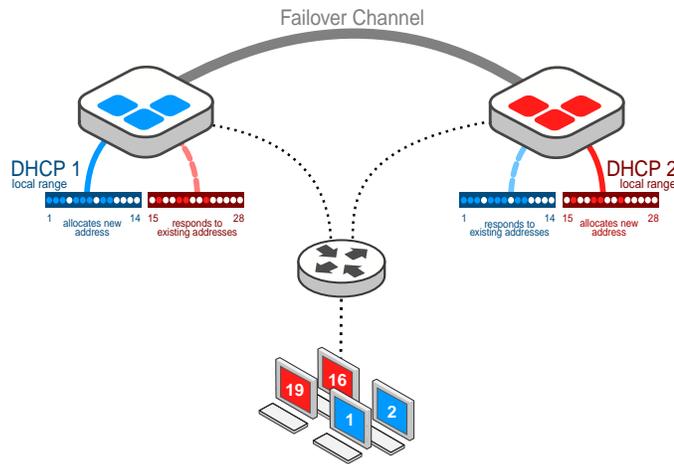


Figure 29.1. DHCP failover operating in normal state

When configuring a One-to-One or One-to-Many smart architecture, the drop-down list *Split leases* allows to choose if the primary server or the secondary server should respond to the DHCP clients requests or choose to balance the responses between the servers. In other words, when configuring your smart architecture, you can set the failover to respect a master/backup configuration or a load balancing configuration.

Which is why in *Normal* state, a server might seem to ignore a client request.

- **Master/Backup Configuration**

If you choose this configuration, you can decide which server of the failover answers to all the requests: either the primary (Prefer master) or the secondary (Prefer backup). For more details, refer to the sections [DHCPv4 One-to-Many Smart Architecture](#) and [DHCPv4 One-to-One Smart Architecture](#).

- **Load Balancing Configuration**

If you choose this configuration, you can balance the responses equally and make both servers respond to the DHCP clients' requests. The standard load balancing algorithm specifies which server answers DHCP requests: this deterministic hash algorithm operates on the clients' information, their MAC address, to equally assign a set of clients to one server and the rest to the other server. The hash is performed on every broadcast message sent out by DHCP clients, it produces a number between 0 and 255 that the servers are able to interpret and divide equally. In addition, when configured in load balancing, EfficientIP DHCP servers are able to detect if a client has not received a response yet from its failover peer. Thanks to the field *secs* of the DHCP client message, the server identifies which clients are making a request for the first time (the field value is zero) or if it is a retry (the field value is nonzero). In the case of a retry, the first available server responds no matter the DHCP client hash number.

Operating in Communications-interrupted State

When operating in *Communications-interrupted* state, each server is operating independently but assumes that its partner is still operating.

The secondary server might be operating and simply unable to communicate with the other server or might not be operating. Each server responds to the full range of DHCP client messages that it receives, but in such a way that graceful reintegration is always possible when its partner comes back and contacts it again.

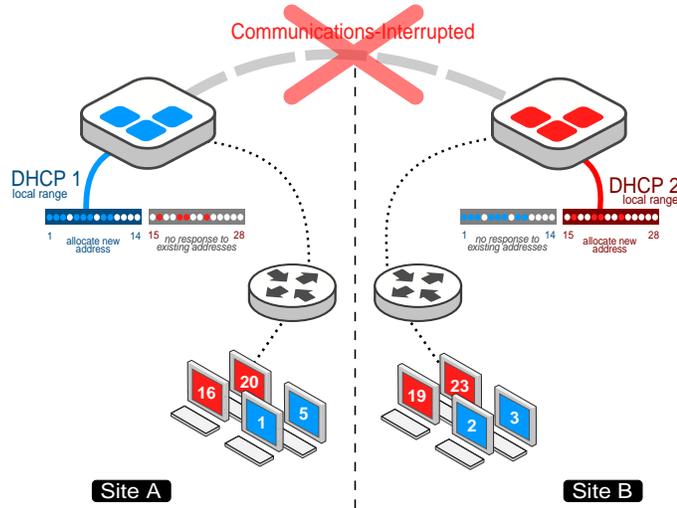


Figure 29.2. DHCP failover operating in communications-interrupted state

Operating in Partner-down State

For a variety of reasons, is it possible that one member of a DHCP failover pair might stop operating. This could be the result of a planned or unplanned outage. In order to provide the best possible service when one member of a failover pair is down, the other can be placed in the *Partner-down* state.

When operating in *Partner-down* state, a server assumes that its partner is not currently operating but does make allowances for the other server's set of DHCP clients as long as the MCLT has not passed. That way, any lease that was allocated by the other server while they were in communication-interrupted state has expired and the remaining server can safely allocate leases to all the DHCP clients of the failover. Once the `MCLT` expires, the server responds to all DHCP client requests, it can reclaim any available IP address that belongs to its peer.

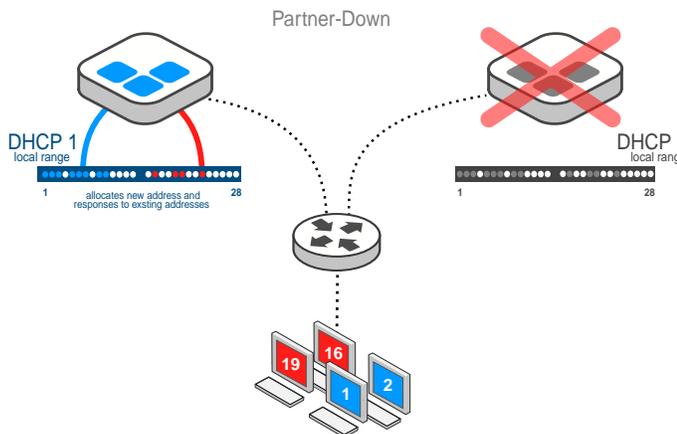


Figure 29.3. DHCP failover operating in partner-down state

Once the peer server comes back up, it automatically connects to its failover channel to change back to the *Normal* operational state. Once again, it has to wait until the [MCLT](#) passes to reclaim its DHCP clients.

You can manually switch a server to *Partner-down*. It allows to better control the DHCP service, for instance before moving a server: the administrator can manually switch the secondary server of a failover channel to *Partner-down*. For more details regarding this option, refer to the section [Switching a DHCP server to Partner-down](#).

In a One-to-One DHCP smart architecture, the administrator can also set an *Automatic switch to partner-down delay (in hours)* after which a server in *Communications-interrupted* state should automatically switch to *Partner-down*. For more details, refer to the section [DHCPv4 One-to-One Smart Architecture](#) of the chapter Managing DHCP Smart Architectures.

Browsing the DHCP Failover Channels Database

To display the list of failover channels

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All failover channels**. The page refreshes.
4. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.

To display a failover channel properties page

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All failover channels**. The page refreshes.
4. On the right-end side of the menu, click on [v4](#) or [v6](#) depending on your needs. The page refreshes and the button turns black.
5. At the end of the line of the failover channel of your choice, click on **■**. The properties page opens.

Getting Familiar with the Failover Channels' Columns

The DHCP module provides a failover channels page both in IPv4 and IPv6.

In contrast with One-to-One and One-to-Many smart architectures that include as many failover channels as physical secondary servers, the failover channel of a Single-Server or a Split-Scope architecture is virtual. It links the managed server(s) to the smart architecture that act as a configuration backup for the `dhcpd.conf` file. For more details, refer to the section [DHCP Failover Principles and Operational States](#).

The concept of failover channels is not very widespread in IPv6. Still, awaiting for its implementation, SOLIDserver already offers a listing page for the virtual failover channels that provides a backup of the smart architectures configuration. For more details, refer to the DHCPv6 architectures of the section [Implementing DHCP Smart Architectures](#).

The DHCPv4 Failover Channels Columns

The DHCPv4 page All failover channels displays 10 columns described in the table below.

The **Split-Scope and Single-Server smart architecture provide few information on this page** as their failover is virtual and therefore cannot be edited. For both architectures, *None* is displayed in every column except *Name*, *Smart DHCP* and *Status*.

Table 29.1. The columns on the page All failover channels in IPv4

Column	Description
Name	The name of each failover channel that you set when creating the smart architecture.
Server	The name of the server or smart architecture.
Type	The failover channel type: either <i>Primary</i> or <i>Secondary</i> .
Local address	The IP address of the primary server, or Master, in the smart architecture.
Local port	The port number dedicated to the failover on the smart architecture primary server.
Remote address	The IP address of the secondary server of the smart architecture.
Remote port	The port number dedicated to the failover on the smart architecture secondary server.
Split	The leases' split configuration between the servers: <i>Balanced</i> , <i>Prefer backup</i> or <i>Prefer master</i> .
State	The failover operational state, either <i>Normal</i> , <i>Startup</i> , <i>Recovering</i> , <i>Partner-Down</i> , <i>Communications-interrupted</i> , <i>Down</i> , <i>Unknown state</i> or <i>N/A</i> . For more details, refer to the table The different failover states below.
Multi-status	Messages regarding the failover channel: emergency, warning, critical, error or informational, if relevant. For more details, refer to the section Understanding the Column Multi-Status .
Status	The failover channel status: either <i>OK</i> , <i>Delayed create</i> or <i>Delayed delete</i> .

The column **State** indicates the failover operational state:

Table 29.2. The different failover states

Failover state	Description
✔ <i>Normal</i>	The server is configured and functions correctly. The failover channel is operational.
⦿ <i>Startup</i>	The failover channel is synchronizing. The failover channel is operational.
⦿ <i>Recovering</i>	The server is recovering from a partner-down state. The failover channel is operational.
⚠ <i>Partner-down</i>	The other server of the failover is <i>Down</i> . The failover channel is operational.
⚠ <i>Communications-interrupted</i>	The server is in communications-interrupted. The failover channel is (not?) operational.
⚠ <i>Down</i>	The server is down. The failover channel is not operational.
⦿ <i>Unknown state</i>	The failover configuration for the smart architecture is incorrect. The failover channel is not operational.
<i>N/A</i>	The failover channel is virtual, therefore returning a state is <i>not applicable</i> , as is the case of Split-Scope and Single-Server smart architectures.

The DHCPv6 Failover Channels Columns

The DHCPv6 page All failover channels displays 9 columns. Considering that the failover channel in IPv6 is basically *virtual* as it is, the *State* column remains empty.

The **Split-Scope and Single-Server smart architecture provide few information on this page**. For both architectures, you can find *N/A* displayed in the port related columns.

Table 29.3. The columns on the page All failover channels in IPv6

Column	Description
Name	The name of each failover channel that you set when creating the smart architecture.
Type	The failover channel type: either <i>Primary</i> or <i>Secondary</i> .
Local address	The IP address of the primary server, or Master, in the smart architecture.
Local port	The port number on the smart architecture primary server dedicated to the failover.
Remote address	The IP address of the secondary server in the smart architecture.
Remote port	The port number of the smart architecture secondary server dedicated to the failover.
State	The connection state between the two servers. As nowadays there is no failover per se in IPv6, this column is empty.
DHCP name	The smart architecture name.
Status	The failover channel status: either <i>OK</i> , <i>Delayed create</i> or <i>Delayed delete</i> .
Multi-status	Messages regarding the failover channel: emergency, warning, critical, error or informational, if relevant. For more details, refer to the section Understanding the Column Multi-Status .

Switching a DHCP server to Partner-down

There are several DHCPv4 operational states in the failover protocol: *Normal*, *Communications-interrupted* and *Partner-down*. When one of the managed servers is unable to communicate with the other, or is down, the failover channel switches to the *Communications-interrupted* state. At that point, you can choose to place the other server in the *Partner-down* state and keep making allowances. There are two ways of switching a server in partner-down: either you automate the switch or you switch the running server manually.

You can automate the switch for servers managed via One-to-One smart architectures, the administrator can also set an *Automatic switch to partner-down delay (in hours)* after which a server in Communications-interrupted state should automatically switch to Partner-down. For more details, refer to the section [DHCPv4 One-to-One Smart Architecture](#) of the chapter Managing DHCP Smart Architectures.

To manually switch a server to partner-down, you can simply break the failover channel following the procedure below, SOLIDserver automatically switches the right server to *Partner-down*. For more details, refer to the section [DHCP Failover Operational States](#).

To manually switch a server to partner-down

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All failover channels**. The page refreshes.
4. On the right-end side of the menu, make sure that the button is black.
5. Tick the failover channel(s) you want to break.

6. In the menu, select  **Edit > Switch to partner-down**. The **Switch to partner-down** wizard opens.
7. Click on  to complete the operation. The report opens and closes. The page **All failover channels** is visible again. In the column **State**, the failover channel has switched to *Partner-down*.

Chapter 30. Configuring DHCP Options

The DHCP dynamically distributes addresses, but also offers the possibility of providing configuration information and other specific controls to the server clients. These pieces of information are called DHCP options.

Most standard DHCP options are currently detailed in the *RFC 2132* recommendation, "DHCP Options and BOOTP Vendor Extensions". Even if most DHCP servers offer several options, the vast majority of DHCP clients are generally conceived to request and take charge of just a sub-part of the ensemble of standard RFC options.

SOLIDserver offers to manage 4 types of DHCP options:

- **Internal options of the DHCP server:** these options allow to configure the global behavior of the DHCP server when it processes DHCP requests. These options do not have DHCP option code number and they are only available on the EfficientIP's DHCP engine provided with SOLIDserver appliances or ISC DHCP software. These options are not sent to the DHCP client. For more details regarding internal server options, refer to the section [Server Parameters](#) in the appendix [MAC Address Types References](#).
- **Client side options:** these options are sent from the DHCP client to the DHCP server to achieve predefined series of actions, for instance *vendor-class* or *hostname* options. If these options can be processed by the server, their content cannot be configured from the server side.
- **Predefined server side options:** these options are predefined and they cannot be modified. Most of these options are common and include options like: *routers*, *domain-name*, *name-server*. These options sent from the server to the client describe network configuration settings and various services available on the network.
- **Custom server side options:** these options can be added and/or modified according to the DHCP clients requirements. These options sent from the server to the client describe network configuration settings and various services available on the network.

SOLIDserver allows you to apply, edit or delete DHCP options.

EfficientIP's DHCP organization allows you to apply DHCP options on three hierarchical levels: the server, the scope and the range. **Microsoft DHCP servers do not allow options configuration on DHCP ranges.**

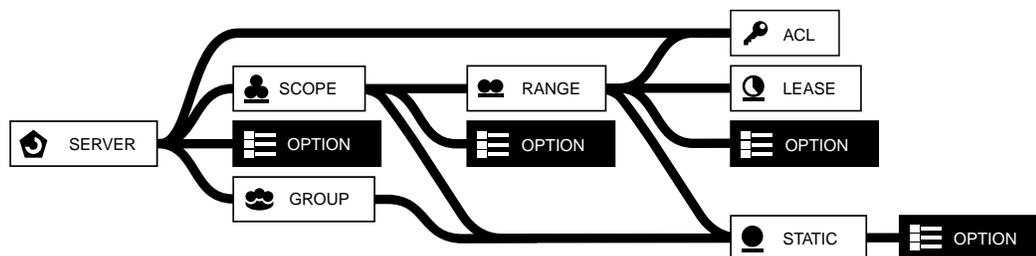


Figure 30.1. The DHCP option in the DHCP hierarchy

In the above configuration plan, the DHCP server options and maximum lease time have been defined to the DHCP server globally; these two options are propagated to the scope and the range.

You might also observe that the default router has been configured both in the DHCP scope and range. Only, in this case, the default router defined by the range is taken into account.

When it comes to DHCPv6, the options configuration of EfficientIP DHCP servers is roughly the same, you can configure options at server and scope level that should propagate to the lower levels. You can also set options at group level or directly on a specific static reservation. However, it is not possible to set DHCP options to a range or a lease.

Setting DHCP Options

DHCP options can be configured from the properties pages of different DHCPv4 objects such as: server, scope, range, static, group, and ACL. In IPv6, the DHCP options can be set at server, scope, static and group level: there are no ACLs and the range DHCP options are not editable.

The **options setting apply to a DHCP client according to a defined precedence**. Options are arranged into a hierarchy in order to respect the following ranking:

- An option set at ACL level overrides all other options.
- An option set at static level overrides options at the following levels: group, range, scope and server.
- An option set at group level overrides options at scope and server level.
- An option set at range level overrides options at scope and server level.
- An option set at scope level overrides options at server level.
- An option set at server level is overridden by all other options.

Options can be indifferently applied to the DHCP objects. However the application of options on this hierarchy depends on the technical constraints of the devices of your network: the devices/clients connected to the network can have an impact on the configuration efficiency.

Basically, the options should be configured by starting from the top of the DHCP tree hierarchy (server) in order not to configure the same options over and over again on each object. Usually options specified at server level are global or applied for a default setup. Everything that was set at server level propagates onto the lower objects, therefore you can configure a common set of options and then add other options to the other objects to match clients needs. If you do not configure the same options repeatedly to several objects, your DHCP configuration is simpler to manage.

The vendors' DHCP servers that SOLIDserver can manage do not share the same internal architecture and cannot be managed in the same way. For instance, contrary to EfficientIP DHCP server, Microsoft DHCP server does not support the configuration of options at range level. Besides, **only the options identified by a number are supported by the Microsoft DHCP service**.

To configure DHCP options in DHCPv4

1. In the sidebar, go to **DHCP > Servers, Groups, Scopes, Ranges** or **Statics** depending on your needs. The page opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. At the end of the line of the object of your choice, click on **⌵**. The properties page opens.
4. In the panel **DHCP options**, click on **EDIT**. The wizard **Configure DHCP options** opens.

5. Select the **Option category** that suits your needs. The available fields refresh.
6. Edit the fields of your choice. For more details regarding options parameters, refer to the part [Customizing DHCP Options](#) below.
7. Click on **OK** to complete the operation. The report opens and closes. The modifications are visible in the panel.

To configure DHCP options in DHCPv6

1. In the sidebar, go to **DHCP > Servers, Scopes or Statics** depending on your needs. The page opens.
2. On the right-end side of the menu, click on **v6**. The page refreshes and the button turns black.
3. At the end of the line of the object of your choice, click on **⚙**. The properties page opens.
4. In the panel **DHCP options**, click on **EDIT**. The wizard **Configure DHCP options** opens.
5. Select the **Option category** that suits your needs. The available fields refresh.
6. Edit the fields of your choice. For more details regarding options parameters, refer to the part [Customizing DHCP Options](#) below.
7. Click on **OK** to complete the operation. The report opens and closes. The modifications are visible in the panel.

Note that you can aggregate range and static options on scopes. For more details, refer to the [Aggregating DHCP Options from Ranges or Statics](#).

Customizing DHCP Options

You can define DHCP custom options for specific DHCP clients like special terminal devices or IP phones. Each value of DHCP option is built by the DHCP server according to a predefined data type, structure of data types or array of types. The graphical user interface allows the administrator of a DHCP server to define the custom data type according to the requirements of the DHCP clients.

To add a custom DHCP option in DHCPv4

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All option definitions**. The page refreshes.
4. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
5. In the menu, click on **+ Add**. The wizard opens.
6. In the list **DHCP server**, select the server on which you want to specify the custom option.
7. Click on **NEXT**. The next page of the wizard opens.
8. In the field **Option name**, name the custom option. The option is named *option <your-option-name>* in the column **Name**.
9. In the field **Option space**, you can fill in the option space parameter that is used to build encapsulated options.

If the space name you chose does not exist, it is created. If you do not specify anything, the default space *dhcp* is used.

10. In the field **Option code**, enter an option code. This code is a number between 1 and 255.

Keep in mind that if you are creating a code within the *dhcp* space, you must define a code greater than 128. The option codes included between 1 and 128 are usually reserved: using a code included in that range of numbers would overwrite existing options.

11. In the drop-down list **Parameter counter**, select the number of parameters you want to set for that option. You can select up to 6 parameters with the corresponding number of fields appearing.
12. In the drop-down list **Parameter <number>**, you have to choose one of the parameters below:

Table 30.1. DHCP options parameter types

Parameter type	Description
IP address	An IPv4 address.
Boolean	A flag accepting a value of either true or false (or yes or no).
Text	An ASCII text string (the same as the text data type) or a list of hexadecimal characters separated by colons. Formatting to distinguish an ASCII text string from a hexadecimal string is important.
8 bits value	A numeric range of the following possible values 8-bit unsigned integer: from 0 to 255 or signed: from -128 to 127.
16 bits value	A numeric range of the following possible values 16-bit signed integer: from -32,768 to 32,767
32 bits value	An ASCII text string (the same as the text data type) or a list of hexadecimal characters separated by colons. Formatting to distinguish an ASCII text string from a hexadecimal string is important. For more details, refer to the section below.
Encapsulate ...	The option parameter <i>Encapsulate <option space></i> is only available for smart architectures managing DHCPv4 servers. It allows to encapsulate options and information, for instance <i>Encapsulate MSFT</i> , <i>Encapsulate MSUCClient</i> , etc. The <i><option space></i> available in the list may vary.

Keep in mind that the encapsulated options' type is binary but equivalent to the text format. Its value is set in hexadecimal and looks as follows: `\x01\xA2\x45\x12`.

If you selected more than one **Parameter counter**, you need to repeat this step for each one them.

13. In the drop-down list **Type is array**, select one of the values below.

Table 30.2. DHCP options array configuration

Parameter type	Description
No	None of the configured options is an array.
Type is array	The last parameter is an array.
Type is global array	Several parameters are arrays.

14. The field **Type** sums up the selected parameters. Each letter that appears in this field corresponds to a parameter. For instance, if you specify an array of IP addresses the type should be **IA**, if you specify an array of repeated addresses plus a boolean the type should be **IfA**.
15. Click on to complete the operation. The report opens and closes. The option is listed.

With DHCPv6, you also have the possibility to add custom options. However, there are fewer parameters available.

To add a custom DHCP option in DHCPv6

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All option definitions**. The page refreshes.
4. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
5. In the menu, click on **+ Add**. The wizard opens.
6. In the list **DHCP server**, select the server on which you want to specify the custom option.
7. Click on **[NEXT]**. The next page of the wizard opens.
8. In the field **Option name**, name the custom option. The option is named "option youroption-name" in the *Name* column.
9. In the field **Option space**, you can fill in the option space parameter that is used to build encapsulated options.

If the space name you chose does not exist, it is created. If you do not specify anything, the default space *dhcp6* is used.

10. In the field **Option code**, enter an option code. This is a number from 1 to 255.

If you are creating a code within the *dhcp* space, you must define a code greater than 128. The option codes included between 1 and 128 are usually reserved: using a code included in that range of numbers would overwrite existing options.

11. In the drop-down list **Parameter counter**, select the number of parameters you want to set for that option. You can select up to 6 parameters with the corresponding number of fields appearing. In each drop-down list, you have to choose one of the parameters below:

Table 30.3. DHCPv6 options parameter types

Parameter type	Description
IP address	An IPv4 address.
Boolean	A flag accepting a value of either true or false (or yes or no).
Text	An ASCII text string (the same as the text data type) or a list of hexadecimal characters separated by colons. Formatting to distinguish an ASCII text string from a hexadecimal string is important.
8 bits value	A numeric range of the following possible values 8-bit unsigned integer: from 0 to 255 or signed: from -128 to 127.
16 bits value	A numeric range of the following possible values 16-bit signed integer: from -32,768 to 32,767
32 bits value	An ASCII text string (the same as the text data type) or a list of hexadecimal characters separated by colons. Formatting to distinguish an ASCII text string from a hexadecimal string is important. For more details, refer to the section below.
Encapsulate server	With DHCPv6, only the Encapsulate server option is available to the servers managed via a smart architecture.

12. For each parameter, one or several boxes are available. Tick the boxes of your choice:

Table 30.4. DHCP options array configuration

Parameter type	Description
No	None of the configured options is an array.

Parameter type	Description
Type is array	This parameter is an array.
Type is global array	Several parameters are arrays.

13. The **Type** field sums up the selected parameters. Each letter that appears in this field corresponds to a parameter. For instance, if you specify an array of IP addresses the type should be **IA**, if you specify an array of repeated addresses plus a boolean the type should be **IfA**.
14. Click on to complete the operation. The report opens and closes. The option is listed.

DHCP Vendor Class Identifier

The vendor class identifier option is used by DHCP clients to specify their vendor type and configuration if need be. The information is a string of n octets, interpreted by servers. Vendors may choose to define specific vendor class identifiers to convey particular configuration or other identification information about a client. For example, the identifier may encode the client's hardware configuration. Servers not equipped to interpret the class-specific information sent by a client must ignore it (although it may be reported). On the contrary, the servers that respond should only use option 43 to return the vendor-specific information to the client.

With DHCPv6, the RFC 3315 defines the Vendor-specific Information Option. SOLIDserver provides it through the *option dhcp6.vendor-opts* (option 17) in the list All option definitions.

Option 82: Relay Agent Information

To put it simply, DHCPv4 Option 82 is the *DHCP Relay Agent Information* option. The DHCP relay agent and Option 82 are defined in RFC 3046. Option 82 was designed to allow a DHCP relay agent to insert circuit specific information into a request that is being forwarded to a DHCP server. Specifically the option works by setting three sub-options: *circuit ID*, *remote ID* and *GIADDR*.

- The field *circuit ID* generally contains information describing the port location that the DHCP request is coming in from. It may contain additional information that helps describe which IP address should be assigned out, such as the VLAN ID, a wireless modem or an ATM virtual circuit. This value must be unique for a particular switch or router that is providing the Relay Agent function. The value must also stay the same if modules are installed or removed in the Switch or Router that implements the Relay Agent. Therefore, having subfields representing the Module, Slot and Port is highly recommended.
- The field *remote ID* is intended to carry information describing the device at the remote end of the link. However, in Ethernet systems, this is typically the MAC address of the Relay Agent. This is not particularly useful since the MAC address would change if the Relay Agent was ever replaced. Building a DHCP server database using the MAC address of the Relay Agent would require that the table be rebuilt every time one of the relay agents was replaced. Some vendors have modified this field to use the IP address of the Relay Agent or some other string describing the relay agent. This field must be unique to the entire network.
- The *GIADDR* (or Gateway Address) field is part of the normal DHCP message. It contains the IP address of the Relay Agent. Since IP addresses must be unique, this field is unique for the entire network.

By combining the *GIADDR* and the *circuit ID*, a network wide unique string can be created. This string can be used for table lookup in the DHCP server. We called this string a pseudo MAC address, since most DHCP servers do a MAC to IP mapping in their databases.

In its default configuration, the *DHCP Relay Agent Information* option passes along port and agent information to SOLIDserver DHCP server. It is useful in statistical analysis, as well as, indicating where an assigned IP address physically connects to the network. It may also be used to make DHCP decisions based on where the request is coming from or even which user is making the request.

The following actions should be performed by the SOLIDserver DHCP when receiving a DHCP-DISCOVER or DHCPREQUEST message with Option 82 set:

1. Relay Agent Information option is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server.
2. Relay Agent Information option is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server.
3. Servers recognizing the Relay Agent Information option may use the information to select the IP address or other parameter assignment policies through the SOLIDserver ACL.
4. Switch or Router (as the DHCP relay agent) intercepting the DHCP requests, appends the *circuit ID* with *remote ID* into the option 82 fields and forwards the request message to SOLIDserver DHCP server.

The following procedure explains how to create an ACL rule allowing to restrict the IPv4 address range to select or to send specific DHCP options according to the option 82 sent to the SOLIDserver DHCP server.

To create an ACL based on the option 82: Circuit ID within the leases user interface

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. On the right-end side of the menu, click on **v4**. The page refreshes and the button turns black.
3. In the column **Name**, click on the server or smart architecture of your choice to display the scopes it contains.
4. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
5. Click on **All ACLs**. The page refreshes.
6. In the menu, click on **+ Add**. The wizard **DHCP ACL parameters** opens.
7. In the field **ACL name**, name your ACL.
8. In the drop-down list **Predefined ACL**, select *None*.
9. In the field **ACL rule**, type in the command below.

```
match if (substring(option agent.remote-id,0,6) = "dslam1");
```

It sets up an ACL that filters the DHCP option 82 as long as the first letters of the client's remote-id match *dslam1*. You can set the keyword of your choice instead.

10. Click on **OK** to complete the operation. The report opens and closes. The ACL is listed.

Once the ACL is created, you can apply it to a DHCPv4 range to allow or restrict the access to all clients that match this ACL rule. ACL can also be used to send specific DHCP options to the clients that match this ACL rule. Edit the properties of the ACL to setup its DHCP option policies.

The Relay Agent Information with DHCPv6

With DHCPv6, the client ID, circuit ID and remote ID are not supported. It is impossible therefore to retrieve these pieces of information separately, much less displaying them in a listing template on the leases page. This information might be delivered by the agent in DHCPv6 but the appliance does not retrieve it at server level.

The equivalent of the option 82 relay agent would be the DHCPv6 option 9 (relay message option) and the option 47 (relay data option).

Option 43: Vendor Specific Information

Option 43 was designed to exchange vendor-specific information between DHCPv4 servers and clients. It was defined in the RFC 2132 as part of the *DHCP Options and BOOTP Vendor Extensions*.

By default, when you add a DHCP smart architecture or an EfficientIP DHCP server, the option 43 is created. This default option cannot be edited.

Within SOLIDserver, the vendor-specific information is stored in an ACL. Any client matching the vendor information is attributed a set of options that you can configure through option definitions. To properly setup option 43 on a DHCPv4 server in the GUI you need to:

1. **Retrieve the vendor-class identifier** from the DHCP handshake.
2. **Create a new ACL** that contains the vendor-class identifier.
3. **Create as many DHCP option definitions as needed** using the ACL as option space.
4. **Configure the server ACL DHCP options** to:
 - a. Set the Vendor option space that triggers the option 43 behavior on all the clients matching the vendor-class identifier.
 - b. Set the value of your choice on all the option definitions you created.

Once the configuration is complete, the clients matching the vendor-class identifier are automatically attributed the option definitions specified.

To retrieve the vendor-class identifier

1. With a packet analyzer, perform a network capture of the DHCP handshake.
2. Open the network capture.
3. In the **Bootstrap Protocol** section, look for the *Vendor class identifier*. It is listed between double quotes among the options, as illustrated below.

```
Bootstrap Protocol
Message type: Boot Request (1)
Hardware type: Ethernet
Hardware address length: 6
Hops: 1
Transaction ID: 0x12adb727
Seconds elapsed: 0
Bootp flags: 0x8000 (Broadcast)
Client IP address: 0.0.0.0 (0.0.0.0)
Your (client) IP address: 0.0.0.0 (0.0.0.0)
Next server To address: 0.0.0.0 (0.0.0.0)
Relay agent IP address: 172.16.206.2 (172.16.206.2)
Client MAC address: Polycom_e5:fa:69 (00:04:f2:e5:fa:69)
Client hardware address padding: 00000000000000000000
Server host name not given
```

```

Boot file name not given
Magic cookie: DHCP
Option: (t=53,l=1) DHCP Message Type = DHCP Discover
Option: (t=57,l=2) Maximum DHCP Message Size = 1456
Option: (t=55,l=20) Parameter Request List
Option: (t=12,l=16) Host Name - "SEP0004f2e5fa69"Option: (t=60,l=14) Vendor class identifier
= "Nortel-223x-A"
Option: (t=61,l=7) Client identifier
End Option
    
```

To create a new ACL that includes the vendor-class identifier

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. In the menu, click on **+ Add**. The wizard opens.
5. In the list **DHCP server**, select the DHCPv4 server or smart architecture of your choice.
6. Click on **NEXT**. The page **DHCP ACL parameters** opens.
7. In the field **ACL name**, name your ACL.
8. In the drop-down list **Predefined ACL**, select *None*.
9. In the field **ACL rule**, type in the command below.

```
match if option vendor-class-identifier = "<%found-value>";
```

To create a DHCP option definition that uses the ACL value as option space

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All option definitions**. The page refreshes.
4. In the menu, click on **+ Add**. The wizard **DHCP Option Definition** opens.
5. In the list **DHCP server**, select the DHCPv4 server for which you configured the ACL.
6. Click on **NEXT**. The page **DHCP option definition** opens.
7. Configure the option. The accepted code, parameter counter, and type should be mentioned in your device documentation.
 - a. In the field **Option name**, name your option.
 - b. In the field **Option space**, type in the ACL name.
 - c. In the field **Option code**, type in a code following your device documentation.
 - d. In the drop-down list **Parameter counter**, select a value following your device documentation.
 - e. In the drop-down list **Parameter 1**, select a value following your device documentation.
 - f. In the drop-down list **Type is array**, select one of the values below.

Table 30.5. DHCP options array configuration

Parameter type	Description
No	None of the configured options is an array.
Type is array	The last parameter is an array.

Parameter type	Description
Type is global array	Several parameters are arrays.

8. Click on **OK** to complete the operation. The report opens and closes. The option is listed as follows: *<option-space-name>.<option-name>*.

Repeat this procedure for as many option definitions as needed: each definition creates a field in the DHCP options configuration wizard which value you can set in the procedure below.

To configure the server with your DHCP option

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **All ACLs**. The page refreshes.
4. Filter the list if need be.
5. At the end of the line of the ACL of your choice, click on **■**. The properties page opens.
6. Click on **☰** to expand all the panels.
7. In the panel **DHCP options**, click on **EDIT**. The wizard **Configure DHCP options** opens.
8. Configure the vendor-specific identifier match:
 - a. In the drop-down list **Option category**, select *Basic*. The wizard refreshes.
 - b. In the drop-down list **Vendor option space**, select your option, it is listed as follows *Vendor <your-option-name>*.
9. Configure the value of your option definitions:
 - a. In the drop-down list **Option category**, select *Vendor <your-option-name>*. The wizard refreshes.
 - b. Fill in all the option definition fields you created. They are all displayed as follows: *<your-option-definition-name> (<your-option-code>)*. The value expected in each field depends on what settings your configured when creating the option definition.
10. Click on **OK** to complete the operation. The report opens and closes. The option is listed in the panel.
 - In the panel **Main properties**, the field **Rule** contains the value of your ACL: the vendor-specific identifier match conditions.
 - In the panel **DHCP options**, you can see:
 1. The field **Vendor option space** that displays your option name.
 2. A field for each of your option definitions named as follows: **<your-option-name>.<your-option-definition-name>** followed by the value you just set in the DHCP option configuration wizard.

Chapter 31. Configuring DHCPv6 Prefix Delegation

DHCPv6 prefix delegation allows to delimit a number of IPv6 addresses, a delegation range, that you distribute using a specific prefix and deliver independently. The Customer Premises Equipment (CPE) can then use it to allocate addresses to their clients. This replaces the need for Network Address Translation (NAT) in an IPv6 network and is widely required when implementing IPv6 network. DHCPv6 prefix delegation is currently detailed in the *RFC 3633* available on IETF website at <https://tools.ietf.org/html/rfc3633>.

Prerequisites

- Defining the delegation range. You must set the start address and end addresses to define the number of IP addresses available for prefix delegation.
- Specifying a shared network that corresponds to one or more scopes. Any scope included in a shared network can use any of the prefix delegations configured.
- Specifying a prefix that sets the size of the IP address segments delegated between the start and end IP addresses.

Specificities

- DHCPv6 prefix delegations are compatible with DHCP relay mechanisms.
- DHCPv6 prefix delegations prevent deleting the scope(s) they are associated with.
- DHCPv6 prefix delegations prevent editing the shared network of the scope(s) they are associated with.
- DHCPv6 prefix delegations prevent editing the type of the smart architecture they are associated with to an architecture that does not support prefix delegation.

Limitations

- As described in RFC 8156, DHCP prefix delegations are not compatible with DHCPv6 failover mechanisms since they are not yet implemented in DHCPd 4.3 nor 4.4. For more details, refer to the RFC available on IETF website at <https://tools.ietf.org/html/rfc8156>.
- You cannot configure IPv6 prefix delegations on Split-scope smart architectures. Note that, on SOLIDserver, it is not possible to display the leases allocated within the delegation.
- You cannot edit DHCPv6 prefix delegations.
- You cannot specify a start address and end address that match existing DHCP ranges on the page *All ranges* of the server.

Browsing the DHCPv6 Prefix Delegations

DHCPv6 prefix delegations are listed on the page *All prefix delegations (v6)*.

To display the list of DHCPv6 prefix delegations

1. In the sidebar, go to  **DHCP > Scopes**. The page **All scopes** opens.

2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the breadcrumb on the right of **All scopes**, click on **»** to display additional pages.
4. Click on **All prefix delegations (v6)**. The page refreshes.

To display a DHCPv6 prefix delegation properties page

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the breadcrumb on the right of **All scopes**, click on **»** to display additional pages.
4. Click on **All prefix delegations (v6)**. The page refreshes.
5. Filter the list if need be.
6. At the end of the line of the prefix delegation of your choice, click on **[E]**. The properties page opens.

Adding DHCPv6 Prefix Delegations

To add a DHCPv6 prefix delegation on a server, you must define an IP address delegation range, a prefix size and a shared network that corresponds to one or more scopes. This way, an equipment asking for a prefix delegation on any scopes belonging to the selected shared network can be delivered one.

Note that, on SOLIDserver, it is not possible to display the leases allocated within the delegation.

Keep in mind that you cannot edit DHCPv6 Prefix delegations.

To add DHCPv6 prefix delegation

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on **[v6]**. The page refreshes and the button turns black.
3. In the breadcrumb on the right of **All scopes**, click on **»** to display additional pages.
4. Click on **All prefix delegations (v6)**. The page refreshes.
5. In the menu, click on **+ Add**. The wizard **Add a prefix delegation** opens.
6. In the drop-down list **DHCP server**, select the DHCP server or smart architecture of your choice.
7. Click on **[NEXT]**. The next page opens.
8. In the field **Start address**, type in the first IPv6 address of the prefix delegation range.
9. In the field **End address**, type in the last IPv6 address of the prefix delegation range.
10. In the field **Prefix**, type in the prefix to delegate. It defines the size of network segments to delegate.

The number of segments depends on the number of IP addresses in the delegation range. The total number of IP addresses contained in all the segments is inferior or equal to the number of IP addresses contained in the range you specified in the fields above.

11. In the field **Shared network**, type in the name of a shared network. The field auto-completes.

Note that for every scope without selecting an existing shared network, one is automatically created anyway. So you may already have shared network available, they are named after the scope *start_address/prefix*.

12. Click on to launch the wizard. The report opens and closes. The prefix delegation is listed.

In the properties page of the scope, the prefix delegation is listed in the panel **Related prefix delegations**.

Deleting DHCPv6 Prefix Delegations

At any time, you can delete a DHCPv6 prefix delegation.

To delete DHCPv6 prefix delegation

1. In the sidebar, go to **DHCP > Scopes**. The page **All scopes** opens.
2. On the right-end side of the menu, click on . The page refreshes and the button turns black.
3. In the breadcrumb on the right of **All scopes**, click on **»** to display additional pages.
4. Click on **All prefix delegations (v6)**. The page refreshes.
5. Tick the prefix delegation(s) of your choice.
6. In the menu, click on **Delete**. The wizard **Delete** opens.
7. Click on to launch the wizard. The report opens and closes. The prefix delegation is no longer listed.

Chapter 32. Monitoring and Reporting DHCP Data

SOLIDserver provides tools dedicated to monitoring DHCP servers and generating reports. Note that these tools only apply to **DHCP objects managing IPv4 addressing**:

- The **alerts** that you can set on the DHCP pages allow to customize your monitoring. For more details, refer to the chapter [Managing Alerts](#).
- A set of **statistics** are available in dedicated panels of the properties page of DHCP servers, as detailed in the section [Monitoring DHCP Servers From their Properties Page](#).
- A set of data sampling **analytics** are available on the page *Analytics*, as detailed in the section [Monitoring DHCP Servers From the Page Analytics](#).
- A set of **rules** allow to monitor your servers, as detailed in the section [Monitoring DHCP Servers Using Rules](#).
- A number of **reports** on IPv4 servers and scopes are available, as detailed in the section [Generating DHCP Reports](#).

Monitoring DHCP Servers From their Properties Page

On the properties page of a physical or smart DHCP server, some panels are dedicated to monitoring queries and changes. The monitoring panels are the following:

- **DHCP server leases <server-name>** displays a lease dedicated chart for physical servers, if relevant.
- **DHCP server statistics <server-name>** displays query dedicated charts for the physical servers.
- **State log** displays the server logs.
- **Audit** displays all the latest changes performed on the server by the user logged in.

Note that the server analytics panel and lease statistics panel display data retrieved using SNMP, therefore, the graphs are empty if the SNMP is not configured properly. To edit the SNMP parameters of an EfficientIP DHCP server, refer to the section [Editing the SNMP Monitoring Parameters of an EfficientIP DHCP Server](#).

Keep in mind that you can zoom in and out of the charts or decide the period and data to display. For more details refer to the section [Charts](#).

To display the lease statistics of a DHCP server or smart architecture

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The properties page open.
3. Open the panel **DHCP server leases <physical-server-name>** using .

Once the panel is open, you can use the buttons under the chart to: move back  or forward  the start time displayed; zoom in  and out ; refresh the data . The drop-down list allows

to restrict or expand the period of time displayed: *Last 3 hours* (selected by default), *Current hour*, *Day*, *Week*, *Month*, *Year*.

To display the statistics of a DHCP server or smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The properties page open.
3. Open the panel **DHCP server statistics <physical-server-name>** using .

Once the panel is open, you can use the buttons under the chart to: move back  or forward  the start time displayed; zoom in  and out ; refresh the data . The drop-down list allows to restrict or expand the period of time displayed: *Last 3 hours* (selected by default), *Current hour*, *Day*, *Week*, *Month*, *Year*.

To display the state log of a DHCP server or smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The properties page open.
3. Open the panel **State log** using . The panel content retrieves the server state in the logs: *OK*, *KO*, *Invalid settings...* and the time and date for each.

To display the audit of a DHCP server or smart architecture

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The properties page open.
3. Open the panel **Audit** using . The panel displays the latest changes in the database.

Each occurrence specifies the *Date* and time it occurred, the *Service* used, the *User* performing the operation and the server basic information: *DHCP name*, *DHCP type* and *Architecture* if relevant.

By default, it lists the changes carried out by the user logged in, but if they belong to a group with access to the changes from all users, the panel displays all the operations ever performed. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).

Monitoring DHCP Servers From the Page Analytics

The data sampling on the page **Analytics** provides DHCP message dedicated *Top 50* lists for the EfficientIP DHCP physical servers you manage.

The analytics functionality is enabled by default and samples the DHCP messages over specific periods of time. By default, it offers 5-minute samples. To set a shorter or larger periodicity, refer to the section [Configuring the Analytics Retrieval](#).

You can set up an alert on the entries displayed. For more details, refer to the chapter [Managing Alerts](#).

Limitations

- The analytics are only available for EfficientIP DHCP servers.
- Only the first 50 entries matching the selected data are listed. Therefore, if during the selected period of time, 100 pieces of information are identical, the GUI only displays the first 50.
- You might slow your appliance down if you edit the purge mechanism to include more lines or keep data longer than the default 30 days.

Accessing the Page Analytics

The page *Analytics* offers dedicated *Top 50* data samples based on the DHCP messages of the physical servers.

To display the page Analytics

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **Analytics**. The page refreshes.

To display the analytics of a DHCP server or smart architecture

1. In the sidebar, go to [DHCP > Servers](#). The page **All servers** opens.
2. In the column **Name**, click on the server or smart architecture of your choice to display the scopes it contains.
3. In the breadcrumb on the right of the server name, click on **»** to display additional pages.
4. Click on **Analytics**. The page refreshes.

Each column provides and compares message information over the specified sample time:

Table 32.1. The columns on the page Analytics

Column	Description
Server	The name of the physical server. Click on a server name to display only the analytics of that server.
Start date	The time and date when the data retrieval started.
End date	The time and date when the data retrieval stopped. The end date respects the number of minutes set in the periodicity.
Period	The periodicity set for the sample of messages. It is set on a physical server properties page. For more details, refer to the section Configuring the Periodicity .
Address	The IP or MAC address of the requesting client, or the IP address of the relay agent. This column is not displayed for <i>Message types</i> .
Message type	The DHCP message type: <i>DISCOVER</i> , <i>OFFER</i> , <i>REQUEST</i> , <i>ACKNOWLEDGE</i> , <i>NOT ACKNOWLEDGE</i> , <i>RELEASE</i> , <i>DECLINE</i> , <i>INFORM</i> or <i>other</i> . This column is only displayed for <i>Message types</i> .
Total messages	The total number of messages in the DHCP logs during the selected period.
Number of hits	The exact number of times a certain message type was requested by the IP or MAC address for the selected period. For instance, the number of <i>DISCOVER</i> messages for a specific IP address during a 5-minute period.
Total ratio	The percentage of <i>Number of hits</i> compared with the <i>Total messages</i> for the selected period. For instance, the percentage that represents the <i>DISCOVER</i> messages of a specific IP address compared with all the messages in the logs during a 5-minute period.

Monitoring and Reporting DHCP Data

Column	Description
Nb of hits / type	The number of hits matching the selected message type during the selected period. For instance, the number of DISCOVER messages during a 5-minute period. This column is not displayed for <i>Message types</i> .
Relative ratio	The percentage of <i>Number of hits</i> compared with the <i>Nb of hits / type</i> during the selected period. For instance, the percentage that represents the DISCOVER messages of a specific IP address compared with all the DISCOVER messages in the logs during a 5-minute period. This column is not displayed for <i>Message types</i> .

Note that the columns **Start date** and **End date** can be filtered using the keyword **last** to display the data retrieved in the last *X* minutes, *X* being the periodicity set for the server. If no data is retrieved in that period, the list is empty.

Displaying the DHCP Analytics

From the page *Analytics*, you can display physical servers data samples retrieved from its messages. It focuses by default on a sample period of 5 minutes to draw *Top 50* comparisons.

To display specific DHCP analytics data

1. In the sidebar, go to [🔍 DHCP > Servers](#). The page **All servers** opens.
2. In the column **Name**, click on the server or smart architecture of your choice to display the scopes it contains.
3. In the breadcrumb on the right of the server name, click on **»** to display additional pages.
4. Click on **Analytics**. The page refreshes.
5. Under the menu, in the drop-down list **Display**, select the data of your choice. The page refreshes, the selected data is displayed.

All available analytics are detailed in the table below.

6. Under the drop-down list, you can tick the box **Automatic refresh** to automatically refresh the data listed every minute. To edit the page refresh frequency, refer to the section [Editing the Automatic Refresh Frequency](#).

No matter the selected Analytics, the values displayed are representative of a specific *Period* of time which you can edit. For more details, refer to the section [Configuring the Periodicity](#).

Table 32.2. DHCP analytics

Statistic	Description
Top 50 DISCOVER (MAC)	The top 50 clients who sent the most DHCPDISCOVER messages during the configured <i>Period</i> , identified using their MAC address.
Top 50 OFFER (MAC)	The top 50 clients who received the most DHCPOFFER messages during the configured <i>Period</i> , identified using their MAC address.
Top 50 OFFER (IP)	The top 50 clients who received the most DHCPOFFER messages during the configured <i>Period</i> , identified using their IP address.
Top 50 REQUEST (MAC)	The top 50 clients who sent the most DHCPREQUEST messages during the configured <i>Period</i> , identified using their MAC address.
Top 50 REQUEST (IP)	The top 50 clients who sent the most DHCPREQUEST messages during the configured <i>Period</i> , identified using their IP address.
Top 50 ACK (MAC)	The top 50 clients who received the most DHCPACK messages (acknowledged) during the configured <i>Period</i> , identified using their MAC address.

Monitoring and Reporting DHCP Data

Statistic	Description
Top 50 ACK (IP)	The top 50 clients who received the most DHCPACK messages (acknowledged) during the configured <i>Period</i> , identified using their IP address.
Top 50 NAK (MAC)	The top 50 clients who received the most DHCPNAK messages (not acknowledged) during the configured <i>Period</i> , identified using their MAC address.
Top 50 NAK (IP)	The top 50 clients who received the most DHCPNAK messages (not acknowledged) during the configured <i>Period</i> , identified using their IP address.
Top 50 RELEASE (MAC)	The top 50 clients who sent the most DHCPRELEASE messages during the configured <i>Period</i> , identified using their MAC address.
Top 50 RELEASE (IP)	The top 50 clients who sent the most DHCPRELEASE messages during the configured <i>Period</i> , identified using their IP address.
Top 50 DECLINE (MAC)	The top 50 clients who sent the most DHCPDECLINE messages during the configured <i>Period</i> , identified using their MAC address.
Top 50 DECLINE (IP)	The top 50 clients who sent the most DHCPDECLINE messages during the configured <i>Period</i> , identified using their IP address.
Top 50 INFORM (IP)	The top 50 clients who sent the most DHCPINFORM messages during the configured <i>Period</i> , identified using their IP address.
Top 50 RELAY	The top 50 most used DHCP relay agents during the configured <i>Period</i> , identified using their IP address.
Top 50 "Peer holds all free leases"	The top 50 DHCP relay agents on which the failover channel refused leases during the configured <i>Period</i> , identified using their IP address.
Top 50 "Unknown network segment"	The top 50 clients who received <i>unknown network segment</i> messages during the configured <i>Period</i> , identified using their IP address.
Top 50 "Unknown subnet"	The top 50 clients who received message <i>unknown subnet</i> messages during the configured <i>Period</i> , identified using their IP address.
Message types	All the messages sent and received by the DHCP server during the configured <i>Period</i> , identified using their type.

Configuring the Analytics Retrieval

You can configure the analytics retrieval according to your needs. You can edit the sampling period, or *periodicity*, the page automatic refresh frequency, the data retrieval frequency and even its purge frequency.

Editing the Automatic Refresh Frequency

On the page **Analytics**, the box **Automatic refresh** allows to automatically refresh the page display every 60 seconds. You can edit this frequency via a registry database key.

To edit the analytics automatic refresh frequency

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *www.dhcp.stat.refresh* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value. The wizard registry database edit a value opens.
5. In the field **Value**, type in the number of seconds of your choice. By default, it is set to *60*.

6. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the new key value is visible.

Configuring the Periodicity

By default, the data sampling compares the DHCP messages over a periodicity of 5 minutes. The sample time specified in the column *Period* on the page *Analytics*.

You can configure a shorter or larger periodicity on each physical server individually.

Note that no matter the periodicity, the data is available on the page at a frequency specified through the rule 381. To edit that rule, refer to the next section.

To edit the analytics periodicity of a DHCP physical server

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **DHCP analytics**. It indicates if the retrieval is enabled and the periodicity.
4. Click on **EDIT**. The wizard **Configure DHCP analytics** opens.
5. In the drop-down list **Periodicity (min.)**, select the period of your choice: 1, 5, 10 or 15 minutes. By default, 5 is selected.
6. Click on **OK** to complete the operation. The page refreshes, the properties page is visible again.

Configuring the DHCP Analytics Retrieval Frequency

The frequency to which the analytics are displayed on the page *Analytics* is set by the rule 381, *Retrieval of the DHCP analytics data*. By default, every 5 minutes it displays the data comparison results for messages sampled during 1, 5, 10 or 15 minutes, depending on the configured periodicity.

No matter the periodicity you set on the physical server, the data is available in the GUI depending on the rule configuration.

To edit the rule 381 that sets the DHCP analytics data retrieval

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #**, type in 381 and hit Enter. The rule is the only one listed.
4. At the end of the line, click on **⌵**. The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a rule** opens.
6. Click on **NEXT**. The page **Rule filters** opens.
7. Edit the drop-down lists **Day(s) of the week**, **Date of the month**, **Month**, **Hour** and/or **Minute** according to your needs. By default, only the drop-down list *Minute* is set to *Every 5 minutes*.
8. Click on **OK** to complete the operation. The page refreshes, the properties page is visible again.

Configuring the DHCP Analytics Purge Frequency

You can configure the purge mechanism of the analytics retrieval. By default, it is based on:

- The data age. The rule 383, *Configuration of the DHCP analytics purge*, deletes data older than 30 days. You can set it to delete data earlier or later.
- A line count. A registry key deletes data if the analytics database exceeds 100,000 lines - each *Top 50* and *Message types* can reach that many lines. You can set a lower or higher threshold.

Both thresholds work together: once the number of days or the number of lines is met, the unwanted data is deleted.

No matter the way you want to purge your database, **keep in mind that if you set very high thresholds, you may slow down your appliance** because the database contains too much information.

To edit the rule 383 that purges the DHCP analytics

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #**, type in 383 and hit Enter. The rule is the only one listed.
4. At the end of the line, click on . The properties page opens.
5. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a rule** opens.
6. Click on **[NEXT]**. The page **Rule filters** opens.
7. Edit the drop-down lists according to your needs. By default the rule is executed daily at 23:30.

Table 32.3. Filters of the rule 383

Column	Default value
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, 23 is selected.
Minute	Select a period of time, minutes-wise. By default, 30 is selected.

8. Click on **[NEXT]**. The page **Rule parameters** opens.
9. In the field **Number of days**, type in the number of days above which you want the logs database to be purged. By default it is set to 30, logs older than thirty days are automatically deleted.
10. Click on **[OK]** to complete the operation. The page refreshes, the properties page is visible again.

To create a threshold to purge DHCP analytics

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.

3. In the menu, click on **+** **Add**. The wizard **Registry database Add an item** opens.
4. In the field **Name**, type in *dhcp.stats.limit*.
5. In the field **Value**, type in the number of lines above which the data is purged. The default value is *100,000*.
6. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the key is listed.

Exporting the Analytics

From the page *Analytics*, you can export the data listed in a CSV, HTML, XML, XLS or PDF file.

Like any other export, you can retrieve the data immediately or schedule it. For more details, refer to the section [Configuring Exports](#).

Disabling the Analytics

At any time, you can stop retrieving the analytics for any physical server.

To disable the analytics retrieval on a DHCP physical server

1. In the sidebar, go to **DHCP > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **DHCP analytics**. It indicates if the retrieval is enabled and its periodicity.
4. Click on **EDIT**. The wizard **Configure DHCP analytics** opens.
5. Untick the box **Enable analytics collection**. The page refreshes, the drop-down list *Periodicity (min.)* is no longer visible.
6. Click on **OK** to complete the operation. The page refreshes, the properties page is visible again. In the panel, the field **Enable analytics collection** is marked **no**.

Monitoring DHCP Servers Using Rules

In order to monitor DHCP servers efficiently, SOLIDserver allows you to add advanced rules to monitor DHCP events: monitor the scope/range usage or set up an alert dedicated to monitoring the server status.

In the following procedures, we are going to configure a monitoring process to add the rules *105* and *082* that respectively *check DHCP scope/range usage* and *send an alert if a DHCP scope is full*.

Before using the monitoring rules, make sure the server is responding.

To add the rule 105 that checks scopes and ranges usage

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the menu, click on **+** **Add**. The wizard **Add a rule** opens.
4. In the drop-down list **Module**, select *DHCP*.

5. In the drop-down list **Event**, select *Execution of a scheduled rule*.
6. In the list **Rule**, select *(105) Check DHCP scope/range usage*.
7. In the field **Rule name**, name the rule. That name is then listed in the column *Instance*.
8. In the field **Comment**, you can type in a comment if you want.
9. Click on . The page **Rule filters** opens.
10. Set the schedule parameters:

Table 32.4. Scheduled rule filters

Field	Description
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, <i>Every hour</i> is selected.
Minute	Select a period of time, minutes-wise. By default, <i>Every minute</i> is selected.

11. Click on . The page **Rule parameters** opens.
12. In the field **Maximum scope usage**, type in the in percent the maximum scope usage of your choice. By default, *90* is typed in.
13. Click on to complete the operation. The report opens and closes. The rule is listed.

Once the rule 105 is added, add the rule 082. Make sure the SNMP service is configured properly. For more details, refer to the section [Managing the SNMP Service](#).

To add the rule 082 that sends an alert when a scope is full

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the menu, click on  **Add**. The wizard **Add a rule** opens.
4. In the drop-down list **Module**, select *DHCP*.
5. In the drop-down list **Event**, select *Event*.
6. In the list **Rule**, select *(082) Send an alert if a DHCP scope is full*.
7. In the field **Rule name**, name the rule. That name is then listed in the column *Instance*.
8. In the field **Comment**, you can type in a comment if need be.
9. Click on . The page **Rule filters** opens.
10. Click on . The page **Rule parameters** opens.
11. To be notified via SNMP trap:
 - a. In the field **IP address of the SNMP trap**, type in the IP address of the appliance that should receive the SNMP trap.
 - b. In the field **SNMP trap community**, type in the community string for this trap. By default it is *ness*.
12. To be notified via email, in the field **Send a mail to**, type in the email address that should receive the notification.

13. Click on to complete the operation. The report opens and closes. The rule is listed.

Thanks to these two rules, if scopes from your DHCP servers exceed the percentage of usage specified in the rule 105, you are automatically notified via email and/or SNMP trap.

Besides, you can display DHCP scope/range usage, in the panel **State log** of the scope properties page.

Generating DHCP Reports

EfficientIP provides DHCP dedicated reports at server and scope level.

Table 32.5. Available DHCP reports

Page	Report
All servers	Clients Most Used OS
	Server Data Exchanges
	Server Options Comparison
	Server Usage Chart
	Most Used Networks
	Server Usage Evolution Charts
All scopes	Clients Most Used OS
	Scopes Options Comparison
	Scopes Summary

For more details regarding the reports and their generation, refer to the section [Managing Reports](#).

Part VII. DNS

The Domain Name System (DNS) is a hierarchical distributed naming system which main function is to convert an IP address into an intelligible domain name (name resolution) or a domain name into an IP address (reverse resolution). The DNS namespace can be seen as a reversed tree of domains managed by zones.

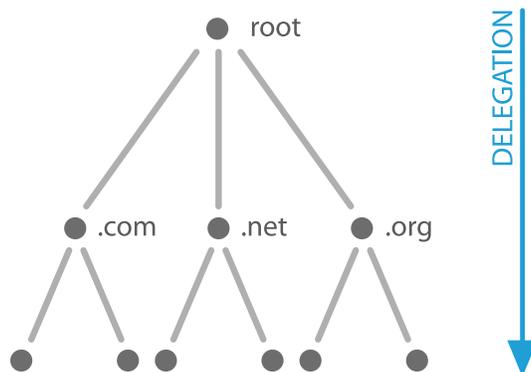


Figure 145. The DNS hierarchy, a reverse tree of delegations

At the root of the structure, the highest zone is represented by a silent dot (.) followed, in order, by the Top-Level Domains (TLDs) and the Second-Level Domains (SLDs). The TLDs are divided into generic TLDs (gTLD) like *.com*, *.org*, *.net* and country code TLDs (ccTLD) like *.us*, *.fr* or *.uk*. The whole access path to a domain reads from right to left: *SLD.TLD*.

At the top of the reverse tree, there are 13 root servers listed alphabetically from A to M and spread out worldwide. They all gather the same information regarding the TLDs and, to avoid being saturated by queries, they delegate names and IP addresses to accredited registrars.

In theory, a client that wants to access a web page would have to follow the whole hierarchy from the root down to the sub-domain. For that reason, DNS servers offer a combination of authoritative, recursive or cache functionalities.

The DNS hierarchy can include up to 4 levels of organization:

- **Server:** the highest level of the hierarchy. It can contain views, zones and records. For more details, refer to the chapters [Managing DNS Servers](#) and [Configuring DNS Servers](#). One or several servers can be managed via a smart architecture to ensure service availability and prevent data or configuration loss. For more details, refer to the chapters [Deploying DNS Smart Architectures](#) and [Managing DNS Smart Architectures](#).
- **View:** an optional level that belongs to a server and contains zones. It allows to grant or limit user access the data of your choice. Depending on who queries a domain, they receive different responses. For more details, refer to the chapters [Managing DNS Views](#) and [Configuring DNS Views](#).
- **Zone:** the second level of the DNS hierarchy. They contain resource records and can be set to resolve names or IP addresses. For more details, refer to the chapters [Managing DNS Zones](#) and [Configuring DNS Zones](#).
- **RR:** the lowest level of the hierarchy. They define zone characteristics. For more details, refer to the chapter [Managing DNS Resource Records](#).

The DNS module also provides:

-
- **DDNS.** You can configure Dynamic Domain Name Server to dynamically take into account in the DNS any IP address assignment updates in the DHCP. For more details, refer to the chapter [Implementing Dynamic Update](#).
 - **RPZ.** You can create and manage Response Policy Zones and set specific rules. For more details, refer to the chapter [DNS Firewall \(RPZ\)](#).
 - **Hybrid DNS.** You can set up a hybrid DNS service that can switch from BIND to NSD or Unbound and vice versa. For more details, refer to the chapter [Hybrid DNS Service](#).
 - **DNSSEC.** You can configure the Domain Name System Security Extensions to use a server as a resolver or sign your zones. For more details, refer to the chapter [DNSSEC](#).
 - **HSM.** You can further secure Master DNSSEC zones with Hardware Security Module encryption. For more details, refer to the chapter [HSM](#).
 - **Monitoring and Reporting.** There are many ways for the reporting and monitoring of DNS servers traffic and activity. For more details, refer to the chapter [Monitoring and Reporting DNS Data](#).
 - **Automation of IPAM and DHCP resources creation.** The advanced properties allow to automate creations in the IPAM or DHCP when you create DNS resources. For more details, refer to the chapter [Managing Advanced Properties](#).

Note that from the module **Dashboards**, you can gather gadgets and charts on *DNS dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 33. Deploying DNS Smart Architectures

The current approach of DNS service management is mainly limited at the single server management level, restricting service configuration and management with a server per server approach even if it is performed from a centralized platform. This approach is insufficient to ensure service reliability, security and easiness of management. It could weaken your DNS architecture because:

- Increases the risk of misconfigurations.
- No Best Practices enforcement to ensure the high security of the network services architecture.
- No automation of architecture deployment and management.
- Difficult and risky architecture changes.

Indeed, even if the configuration has been simplified with the GUI, it is still complex, expensive and requires experts to deploy and configure all servers in coherent architectures of DNS-DHCP services. The smart architecture is a new approach to DNS services management to drastically simplify deployment and administration of your network service. Thanks to the smart architecture, SOLIDserver offers the capability of managing your DNS services not only at server level but at the architecture level.

The smart architecture offers a library of DNS architectures that are ready to apply on a set of servers. **All the DNS smart architectures designed for more than one server can contain several Master servers.** This sets up an even more secure environment: if one Master server crashes or stops responding, the other one takes over and ensures service availability.

The DNS smart architecture library includes:

- Master/Slave.
- Multi-Master.
- Stealth.
- Single-Server.
- Farm.

Smart architecture supports EfficientIP SOLIDserver servers and legacy DNS servers such as:

- Microsoft Windows Server DNS.
- ISC BIND9.
- Nominum ANS.

Smart architecture allows managing other DNS servers supporting DDNS (RFC2136) with the single ability of updating the domains and not the server configuration or the zone configuration. In that way, the server configuration and the zone configuration must be done locally on the server. This configuration is useful when you are only allowed to update zones on a DNS partner.

Master/Slave Smart Architecture

Master/Slave DNS architecture is widely used on the Internet. SOLIDserver supports the Master/Slave DNS architecture within as a smart architecture. A master DNS configuration contains

one or more zones files for which this DNS server is authoritative. The term master is related to the location of the zone data rather than any other operational characteristics. A master is requested to transfer zone data to one or more slave servers whenever the zone file change. The master DNS obtains the zone data locally as opposed to a slave DNS, which obtains its zone data via a zone transfer operation from the master DNS.

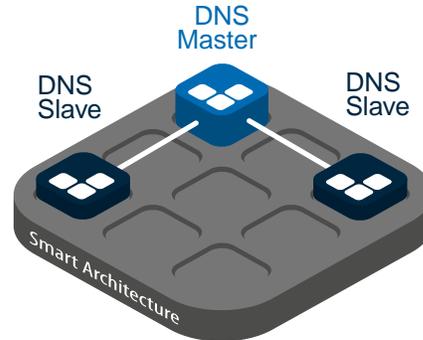


Figure 33.1. DNS Master/Slave smart architecture

Multi-Master Smart Architecture

The DNS Multi-Master architecture is usually selected to allow updates on all servers. Multi-Master smart architecture supports all DNS servers, including Microsoft DNS servers integrated or not in Active Directory, EfficientIP DNS, BIND servers, Nominum ANS engines or all DNS servers supporting DDNS. It can even reproduce Microsoft's Multi-Master behavior.

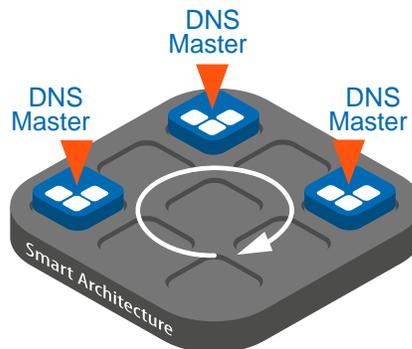


Figure 33.2. DNS Multi-Master smart architecture

With the smart architecture, updating a DNS server can be done from the management console, from a DHCP allocation or from Microsoft DNS clients that update themselves their names by using the Dynamic DNS (DDNS) mechanism:

- When a Multi-Master smart architecture is updated from the management console, the configuration is automatically pushed toward all the DNS servers belonging to the smart architecture.
- When a DNS server receives a dynamic update from a DNS client, the Multi-Master smart architecture replicates the update to all the DNS servers it manages. This replication is automatic and does not require any manual operations.
- When a DHCP server offers a new IP address, the SOLIDserver IPAM appliance updates the Multi-Master smart architecture and, consequently, all the DNS servers it manages.

A primary DNS server is eliminated as a single point of failure. Traditional DNS replication is single-master; it relies on a primary DNS server to update all the secondary servers. Unlike traditional DNS replication, Directory Server Replication is Multi-Master. Changes made to a zone can be replicated to one or more Directory Servers. Which is why we recommend that you refer to your vendor information regarding the Directory Server used and its replication capabilities.

Stealth Smart Architecture

A Stealth DNS architecture is a set of visible DNS servers and a stealth DNS server. A stealth DNS server is defined as a name server that does not appear in the list of the visible DNS servers, which means that its NS resource record is not published among the zone and it does not answer queries from DNS clients and other name servers. Stealth architectures are used in contexts that are sometimes called demilitarized zone (DMZ) or Split servers, and can be defined as having the following characteristics:

- Your organization needs to deploy DNS servers on the Internet.
- Your organization does not want the world to see any of its internal hosts either by interrogation (query or zone transfer) or in the event the DNS service or external servers are compromised.

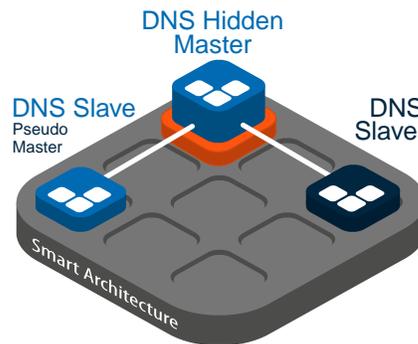


Figure 33.3. DNS Stealth smart architecture

The visible secondary DNS server contains only slave zones, then it is less exposed to DNS attacks because the real authoritative primary server is hidden. Zone transfers can be allowed from the secondary servers as required but they do not transfer or accept transfers from the stealth server.

One of the main advantages of this architecture is that the primary server can be offline for maintenance without causing any interruption to DNS service within the expiration duration (30 days) set for the validity of its zone data.

Single-Server Smart Architecture

A Single-Server architecture manages one DNS server. This allows to keep the server configuration and data in case anything were to happen to the physical server: SOLIDserver would save the configuration and apply it to the next server you add to the architecture.

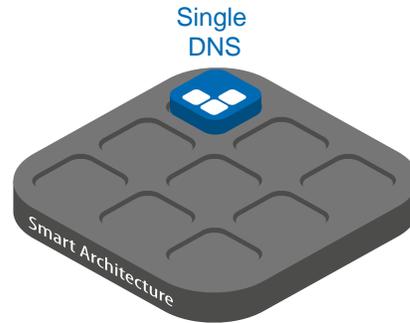


Figure 33.4. DNS Single-Server smart architecture

This architecture is therefore a backup in itself. Moreover, managing a physical server through a Single server architecture eases up any migration or change of architecture. If after a few weeks, for instance, you want to set up a Master/Slave architecture, you can edit the smart architecture, change it to Master/Slave, add another physical server and define which one acts as a master and which one as a slave.

Farm Smart Architecture

The Farm architecture allows to control the DNS service through one or several load balancers. The load balancer receives the DNS clients requests and redirects each query to the least used DNS server at the time of the request. That way, the DNS load is balanced and the service availability is heightened. The load balancer sends the DNS queries to a set of known DNS servers that send back the information needed. The organization of the DNS servers in a Farm architecture is based on the principle of the Master/Slave architecture with one Master server and as many Slave servers as needed.

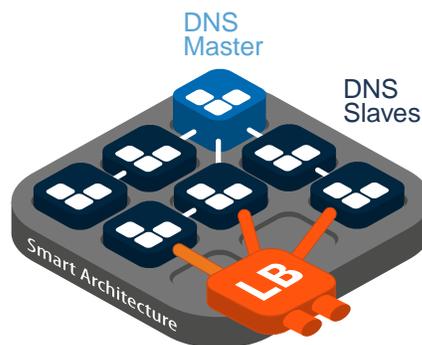


Figure 33.5. DNS Farm smart architecture

The Farm architecture is especially useful for huge configurations where the use of load balancers is necessary.

Chapter 34. Managing DNS Smart Architectures

Once you chose the smart architecture(s) that suit your needs in the chapter [Deploying DNS Smart Architectures](#), you can manage them following the sections below.

Browsing DNS Smart Architectures

Smart architectures are managed from the page **All servers**, listed like physical servers and preceded by the icon . For more details, refer to the section [Browsing DNS Servers](#).

Browsing the DNS Smart Architectures Database

To display the list of DNS smart architectures

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. To display the physical server(s) managed by each smart architecture, click on .

In the column **Name**, all the smart architectures are preceded by the icon . They are listed with the physical servers.

To display a DNS smart architecture properties page

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. At the end of the line of the smart architecture of your choice, click on . The properties pages opens.

Understanding the Smart Architecture Statuses

The column **Status** provides information regarding the smart architectures' configuration.

Table 34.1. Smart architecture statuses

Column	Description
 <i>OK</i>	The smart architecture is operational.
 <i>Invalid settings</i>	The smart architecture does not contain any physical server or is missing one or several.
 <i>Locked synchronization</i>	The server configuration is not viable. This status can be displayed on EfficientIP DNS and EfficientIP DNS Package servers after importing a BIND archive file not properly formed or containing non-supported BIND options. For more details, refer to the section Handling the Status Locked Synchronization .

In addition, the column **Sync** provides additional information regarding the exchanges, synchronization, between the smart architecture and the physical server(s).

Table 34.2. Smart architecture synchronization statuses

Column	Description
 <i>Synchronized</i>	The smart architecture has successfully synchronized the server(s) it manages.

Column	Description
 <i>Busy</i>	The smart architecture is synchronizing the server(s).
 <i>Locked synchronization</i>	The synchronizing failed as the server configuration is not viable: the smart architecture cannot send the configuration file to the physical server(s). For more details, refer to the section Handling the Status Locked Synchronization .

Adding a DNS Smart Architecture

A smart architecture can be configured without DNS servers. It allows you to create the architecture that suits your needs before applying it to one or more DNS servers. It also provides a backup of the management configuration of the server it manages. If your DNS server crashes, you delete it and add a new one on which you apply the same architecture, SOLIDserver remembers the former server's configuration and apply it to the new one.

There are five different kinds of smart architectures: Master/Slave, Multi Master, Stealth, Farm and Single-Server. Keep in mind that every DNS smart architecture sets up an active/active configuration. **The smart architectures Farm, Master/Slave, Multi-Master and Stealth can manage several Master servers.** This sets up an even more secure environment: if one Master server crashes or stops responding, the other one takes over and ensures service availability.

Once the configuration is completed, the DNS smart architecture is listed as a real server on the page *All servers*, the column *Type* indicates the kind of smart architecture.

In the procedures below, we are going to describe the configuration of the DNS smart architectures with the DNS servers they manage, but you can go through the configuration without adding any server and do it later. For more details, refer to the part [Adding a DNS Server into a Smart Architecture](#).

Master/Slave Smart Architecture

The [Master/Slave Smart Architecture](#) is designed to manage at least 2 DNS servers with one DNS server as master and the other(s) as slave (i.e. backup).

To configure a DNS Master/Slave smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 34.3. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration

Field	Description
	allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DNS smart architecture**, select **Master/Slave**.

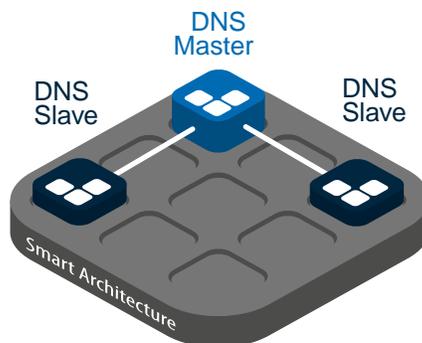


Figure 34.1. DNS Master/Slave smart architecture

7. Click on **NEXT**. The page **DNS servers role configuration** opens.
8. You can select the DNS servers that you want to manage through the smart architecture:
 - a. In the drop-down list **Available DNS servers**, select the master server and click on **+ MASTER**. The server is moved to the **Master DNS server(s) list**. You can add several master servers if you want, in which case if one crashes the other takes over. To remove a server from the list, select it and click on **-**.
 - b. In the drop-down list **Available DNS servers**, select a slave server and click on **+ SLAVE**. The server is moved to the **Slave DNS servers list**. Repeat this action for as many slave servers as needed. To remove a server from the list, select it and click on **-**.

If you do not want to publish one or several name servers/load balancers or enable HSM for this architecture, go to step 9.

9. If you want to publish one or several name servers/load balancers or enable HSM for this architecture, tick the **Expert mode** box. The page reloads.

- a. Click on **NEXT**. The page **Advanced settings** appears.
- b. In the field **NS record**, type in the name server of your choice. It can also be the hostname of an external load balancer.
- c. Click on **ADD**. The name server is moved to the **Published name servers list**. Each record listed is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many NS records as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

- d. You can tick the box **Force Hybrid DNS compatibility** if you intend to manage BIND servers that you might switch to Hybrid in the future. For more details, refer to the chapter [Hybrid DNS Service](#).
 - e. Click on **NEXT**. The page last page of the wizard appears.
 - f. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM. For more details, refer to the chapter [HSM](#).
10. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DNS server and marked **Smart (master/slave)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

During the first addition of a DNS smart architecture, the option allow-transfer is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use as it is inherited by the server's zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Multi-Master Smart Architecture

The [Multi-Master Smart Architecture](#) is designed to manage at least 2 DNS servers: both of them being Masters, there is no Slave server in this configuration. From the management console, a DNS client or a DNS server automatically replicates and updates data on all the DNS servers within this architecture.

To configure a DNS Multi-Master smart architecture

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 34.4. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DNS smart architecture**, select **Multi-Master**.

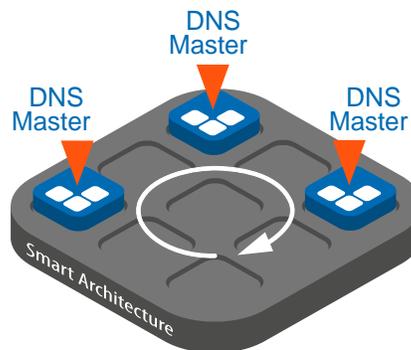


Figure 34.2. DNS Multi-Master smart architecture

7. Click on **NEXT**. The page **DNS servers role configuration** opens.
8. You can select physical servers to manage through the smart architecture:
 - a. In the drop-down list **Available DNS servers**, select a server.
 - b. Click on **+ MASTER**. The server is moved to the **Master DNS servers list**. You can add several master servers if you want, in which case if one crashes the other takes over. To remove a server from the list, select it and click on **-**.

If you do not want to configure any name server or load balancer for this architecture, go to step 9.

9. If you want to publish one or several name servers or load balancers for this architecture, tick the **Expert mode** box. The page reloads.
 - a. Click on **NEXT**. The page **Advanced settings** appears.
 - b. In the field **NS record**, type in the name server of your choice. It can also be the hostname of an external load balancer.
 - c. Click on **ADD**. The name server is moved to the **Published name servers list**. Each record listed is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many NS records as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.
 - d. The field **Compatible with a Hybrid DNS Engine** is marked **Yes**.
 - e. You can tick the box **Force Hybrid DNS compatibility** if you intend to manage BIND servers that you might switch to Hybrid in the future. For more details, refer to the chapter [Hybrid DNS Service](#).
 - f. Click on **NEXT**. The page last page of the wizard appears.
 - g. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM. For more details, refer to the chapter [HSM](#).
10. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DNS server and marked **Smart (multi-master)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

During the first addition of a DNS smart architecture, the option allow-transfer is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use as it is inherited by the server's zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Stealth Smart Architecture

The [Stealth Smart Architecture](#) is designed to manage at least 3 DNS servers: a true Master server hidden from the world, a visible Master server used as decoy and Slave server(s) that do not transfer or accept transfers from the hidden Master server. The Master server can be offline for maintenance without causing any interruption to DNS service within the expiration duration (30 days) set for the validity of its zone data.

To configure a DNS Stealth smart architecture

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields according to the table below:

Table 34.5. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The next page of the wizard opens.
- In the list **DNS smart architecture**, select **Stealth**.

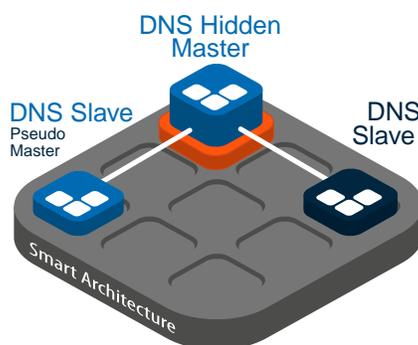


Figure 34.3. DNS Stealth smart architecture

- Click on **NEXT**. The next page of the wizard opens.
- You can select the DNS servers that you want to manage through the smart architecture:
 - In the drop-down list **Available DNS servers**, select the master server and click on **HIDDEN-MASTER**. The server is moved to the **Hidden-master DNS server(s) list**. Repeat this action for as many master servers as needed. To remove a server from the list, select it and click on **REMOVE**.

- b. In the drop-down list **Available DNS servers**, select the slave server you want to use as pseudo master and click on **+ PSEUDO-MASTER**. The server is moved to the field **Pseudo-master DNS server (slave server used as decoy)**. To remove the server from the field, click on **■**.
- c. In the drop-down list **Available DNS servers**, select a slave server and click on **+ SLAVE**. The server is moved to the **Slave DNS servers list**. Repeat this action for as many slave servers as needed. To remove a server from the list, select it and click on **■**.

If you do not want to configure any name server or load balancer for this architecture, go to step 9.

9. If you want to publish one or several name servers or load balancers for this architecture, tick the **Expert mode** box. The page reloads.
 - a. Click on **NEXT**. The page **Advanced settings** appears.
 - b. In the field **NS record**, type in the name server of your choice. It can also be the hostname of an external load balancer.
 - c. Click on **ADD**. The name server is moved to the **Published name servers list**. Each record listed is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many NS records as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

- d. The field **Compatible with a Hybrid DNS Engine** is marked **Yes**.
 - e. You can tick the box **Force Hybrid DNS compatibility** if you intend to manage BIND servers that you might switch to Hybrid in the future. For more details, refer to the chapter [Hybrid DNS Service](#).
 - f. Click on **NEXT**. The page last page of the wizard appears.
 - g. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM. For more details, refer to the chapter [HSM](#).
10. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DNS server and marked **Smart (stealth)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

During the first addition of a DNS smart architecture, the option allow-transfer is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use as it is inherited by the server's zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Single-Server Smart Architecture

The [Single-Server Smart Architecture](#) is designed to manage only one DNS physical server.

To configure a DNS Single-Server smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.

2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 34.6. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DNS smart architecture**, select **Single-Server**.

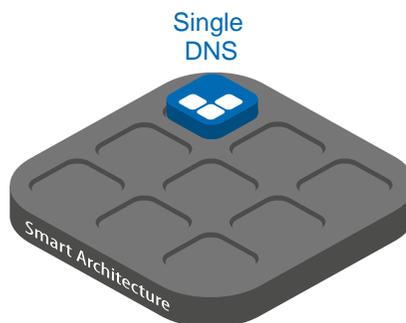


Figure 34.4. DNS Single-Server smart architecture

7. Click on **NEXT**. The last page of the wizard opens.
8. You can select the DNS server that you want to manage through the smart architecture:

- a. In the drop-down list **Available DNS servers**, select the server.
- b. Click on . The server is moved to the **Master DNS servers list**. To remove a server from the list, click on .

If you do not want to configure any name server or load balancer for this architecture, go to step 9.

9. If you want to publish one or several name servers or load balancers for this architecture, tick the **Expert mode** box. The page reloads.
 - a. Click on . The page **Advanced settings** appears.
 - b. In the field **NS record**, type in the name server of your choice. It can also be the hostname of an external load balancer.
 - c. Click on . The name server is moved to the **Published name servers list**. Each record listed is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many NS records as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on or click on to remove it from the list. If you made changes that you do not want to save, click on .

- d. The field **Compatible with a Hybrid DNS Engine** is marked **Yes**.
 - e. You can tick the box **Force Hybrid DNS compatibility** if you intend to manage BIND servers that you might switch to Hybrid in the future. For more details, refer to the chapter [Hybrid DNS Service](#).
 - f. Click on . The page last page of the wizard appears.
 - g. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM. For more details, refer to the chapter [HSM](#).
10. Click on to complete the operation. The report opens and closes. The smart architecture is listed as a DNS server and marked **Smart (single-server)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button on the right-end side of the menu.

During the first addition of a DNS smart architecture, the option allow-transfer is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use as it is inherited by the server's zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Farm Smart Architecture

The [Farm Smart Architecture](#) is essentially a Master/Slave architecture that allows to have a set of master and slave servers accessible through one or several external load balancers that redirect the clients toward the least used server and avoid overloading the service.

To configure a DNS Farm smart architecture

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.

3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the fields according to the table below:

Table 34.7. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

5. Click on **NEXT**. The next page of the wizard opens.
6. In the list **DNS smart architecture**, select **Farm**.

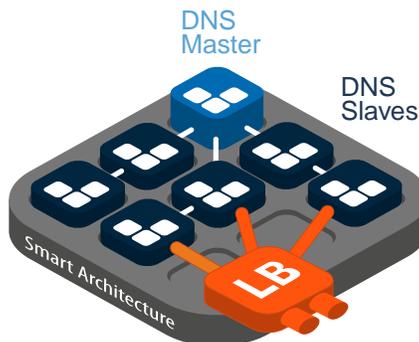


Figure 34.5. DNS Farm smart architecture

7. Click on **NEXT**. The page **DNS servers role configuration** opens.
8. You can select the DNS servers that you want to manage through the smart architecture:

- a. In the drop-down list **Available DNS servers**, select the master server and click on **MASTER**. The server is moved to the **Master DNS servers list**. You can add several master servers if you want, in which case if one crashes the other takes over. To remove a server from the list, select it and click on .
 - b. In the drop-down list **Available DNS servers**, select a slave server and click on **+ SLAVE**. The server is moved to the **Slave DNS servers list**. Repeat this action for as many slave servers as needed. To remove a server from the list, select it and click on .
9. Click on **NEXT**. The page **Advanced settings** opens.
 10. Finish the Farm configuration.
 - a. In the field **NS record**, type in the hostname of your external load balancer if need be. It can also be a name server.
 - b. Click on **ADD**. The name is moved to the **Published name servers list**. Each record is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many load balancers or NS records as needed.

From now on, the DNS clients send their request to one or more load balancers that redirect the requests to the least used server. Note that to run properly, your load balancer must be configured to list all the DNS servers managed by the smart architecture and should be manually updated if you change the list of physical servers managed by the architecture.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

11. If you want to display the Hybrid dedicated fields, tick the box **Export mode**.
 - a. The field **Compatible with a Hybrid DNS Engine** is marked **Yes**.
 - b. You can tick the box **Force Hybrid DNS compatibility** if you intend to manage BIND servers that you might switch to Hybrid in the future. For more details, refer to the chapter [Hybrid DNS Service](#).
 - c. Click on **NEXT**. The page last page of the wizard appears.
 - d. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM. For more details, refer to the chapter [HSM](#).
12. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed as a DNS server and marked **Smart (farm)** in the column **Type**. You can display or hide the physical servers managed through your smart architecture using the button on the right-end side of the menu.

During the first addition of a DNS smart architecture, the option allow-transfer is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use as it is inherited by the server's zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Editing a DNS Smart Architecture

Once created, you can edit a smart architecture to change the servers it manages, edit the server roles, change the smart architecture type, convert a server into a smart architecture or enable HSM.

Adding a DNS Server into a Smart Architecture

Once smart architecture is properly configured and applied, you can add DNS servers whenever you want. First add a DNS server following the section [Managing DNS Servers](#). Make sure you added all the necessary physical servers into the smart architecture because, depending on the DNS smart architecture you chose, if you do not complete the configuration the smart architecture may not run properly.

When you add one or more DNS servers into a smart architecture, the smart data is automatically replicated from the architecture to the DNS servers it manages. So if the smart architecture is empty (first use), the configuration file of the physical DNS server it manages is completely overwritten.

To add a DNS server into a smart architecture

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on **⊞**. The properties page opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
5. In the list **DNS server type**, make sure *DNS smart architecture* is selected. Click on **NEXT**. The next page of the wizard opens.
6. If need be, modify the smart architecture basic parameters according to the table below:

Table 34.8. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .

Field	Description
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- In the list **DNS smart architecture**, modify the type of DNS smart architecture if need be. Click on **NEXT**. The page **DNS servers role configuration** opens.
- In the drop-down list **Available DNS servers**, select the DNS server of your choice.
- Define the role of the server using the buttons **+ HIDDEN MASTER**, **+ PSEUDO MASTER**, **+ MASTER** or **+ SLAVE** depending on the smart architecture. The selected server is moved to the corresponding **Hidden-master**, **Pseudo-master**, **Master** or **Slave DNS servers list**. You can remove the server from the list using **-**. Repeat these actions for as many servers as needed.
- If you are editing a Farm architecture or if you configured NS records on another architecture, click on **NEXT**. The page **Advanced settings** opens. For more details regarding this page, refer to the last steps of the relevant smart architecture addition procedure in the section [Adding a DNS Smart Architecture](#).
- Click on **OK** to complete the operation. The report opens and closes. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu.

Removing a DNS Server from a Smart Architecture

Whenever you want to, you can remove one or more DNS servers from a DNS smart architecture. When you remove one, the configuration applied on this server is conserved on the previously removed DNS server.

To remove a DNS server from a smart architecture

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the smart architecture of your choice, click on **-**. The properties page opens.
- In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
- If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Click on **NEXT**. The next page of the wizard opens.
- Click on **NEXT**. The next page of the wizard opens.
- Select the server to remove and click on **-**. The server is moved back to the drop-down list **Available DNS servers**. Repeat this action for the other servers you want to remove.
- If you are editing a Farm architecture or if you configured NS records on another architecture, click on **NEXT**. The page **Advanced settings** opens. For more details regarding this page,

refer to the last steps of the relevant smart architecture addition procedure in the section [Adding a DNS Smart Architecture](#).

- Click on **OK** to complete the operation. The report opens and closes. If your smart architecture is still managing physical servers, you can display or hide them using the button  on the right-end side of the menu.

Changing the DNS Servers Role within a Smart Architecture

You can easily modify the role of the DNS servers managed by any smart architecture. For instance, you can change a master server into a slave server within a Master-Slave smart architecture at any given time.

To change the role of a DNS server within a smart architecture

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the smart architecture of your choice, click on . The properties page opens.
- In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
- If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Click on **NEXT**. The next page of the wizard opens.
- Click on **NEXT**.
- Click on **NEXT**. The page **DNS servers role configuration** of the wizard opens.
- Select the server you want to modify and click on  next to the corresponding list. The server is moved back to the drop-down list **Available DNS servers**. Repeat this action for any server whose role you want to change.
- In the drop-down list **Available DNS servers**, select the DNS server of your choice.
- Define the role of the server using the buttons **+ HIDDEN MASTER**, **+ PSEUDO MASTER**, **+ MASTER** or **+ SLAVE** depending on the smart architecture. The selected server is moved to the corresponding list. Repeat these actions for the other servers.
- If you are editing a Farm architecture or if you configured NS records on another architecture, click on **NEXT**. The page **Advanced settings** opens. For more details regarding this page, refer to the last steps of the relevant smart architecture addition procedure in the section [Adding a DNS Smart Architecture](#).
- Click on **OK** to complete the operation. The report opens and closes. You can display or hide the physical servers managed through your smart architecture using the button  on the right-end side of the menu. The column **Role** displays the server(s) new role.

Changing the Type of a DNS Smart Architecture

The type of a DNS smart architecture can be easily changed while keeping all DNS configuration and data you already set. For instance, you already have a DNS smart architecture configured in Master-Slave that includes two DNS servers -one in master and the other in slave- and you

plan to change your DNS configuration type from Master-Slave to Multi-Master. By editing the smart architecture, you can change its type and configure the role of servers.

To edit a DNS Smart Architecture type

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on **⊞**. The properties page opens.
3. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DNS server** opens.
4. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. In the list **DNS server type**, make sure *DNS smart architecture* is selected. Click on **[NEXT]**. The next page of the wizard opens.
6. If need be, modify the smart architecture basic parameters according to the table below:

Table 34.9. DNS smart architecture basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Isolated	Tick the box if you do not want your server configuration to update any other module (IPAM or DHCP). It's mainly useful when dealing with migrations. Keep in mind that the smart architecture still receives data if your network configuration allows it. If you tick the box, any behavior set via the drop-down list <i>Advanced properties</i> has to be applied to the smart architecture later on. Before unticking the box make sure that the configuration you set suits your needs.
Use DNS as DNSSEC resolver	Tick the box to activate DNSSEC validation. If you activate the DNSSEC parameters on a smart architecture, all the servers it manages become DNSSEC compliant. For more details, refer to the chapter DNSSEC .
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

7. In the list **DNS smart architecture**, modify the type of your DNS smart architecture. Click on **[NEXT]**. The page **DNS servers role configuration** opens.
8. Select the server you want to modify and click on **⊞** next to the corresponding list. The server is moved back to the drop-down list **Available DNS servers**. Repeat this action for any server whose role you want to change.
9. In the drop-down list **Available DNS servers**, select the DNS server of your choice.

10. Define the role of the server using the buttons `+ HIDDEN MASTER`, `+ PSEUDO MASTER`, `+ MASTER` or `+ SLAVE` depending on the smart architecture. The selected server is moved to the corresponding list. Repeat these actions for the other servers.

If you selected the Master/Slave, Multi-Master, Stealth or Single-Server architecture, go to step 11.

11. If you selected the Farm architecture, click on `NEXT`. The page **Advanced settings** opens.
 - a. In the field **NS record**, type in the hostname of your external load balancer if need be. It can also be a name server.
 - b. Click on `ADD`. The name is moved to the **Published name servers list**. Each record is saved for each zone and displayed on the page *All RRs* of the physical servers managed by the smart architecture.

Repeat these actions for as many load balancers or NS records as needed.

From now on, the DNS clients send their request to one or more load balancers that redirect the requests to the least used server. Note that to run properly, your load balancer must be configured to list all the DNS servers managed by the smart architecture and should be manually updated if you change the list of physical servers managed by the architecture.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on `UPDATE` or click on `DELETE` to remove it from the list. If you made changes that you do not want to save, click on `CANCEL`.

12. Click on `OK` to complete the operation. The report opens and closes. The page **All servers** is visible again. The column **Type** displays your changes.

Converting a DNS Server into a Smart Architecture

To keep a server configuration and avoid configuring a smart architecture to match the server settings before adding it into the smart, you can convert DNS servers into smart architectures.

Keep in mind that once you converted a DNS server into a smart, it is no longer listed on the page *All servers*. You have to add it again to be able to manage it, on its own or from a smart architecture.

During the conversion, you can add DNS servers into the smart architecture. Considering that you might want to manage the server you converted from the smart architecture, we recommend converting the server and then editing the smart to add the servers as detailed in the section [Adding a DNS Server into a Smart Architecture](#).

To edit a DNS Smart Architecture type

1. In the sidebar, go to `DNS > Servers`. The page **All servers** opens.
2. Right-click over the **Name** of the server of your choice. The contextual menu appears.
3. Click on `Edit`. The wizard **Edit a DNS server** opens.
4. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on `NEXT`. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. In the list **DNS server type**, select *DNS smart architecture*.
6. Click on **NEXT**. The next page of the wizard opens, it displays the server details.
7. Click on **NEXT**. The next page of the wizard opens.
8. In the field **DNS smart architecture**, select the architecture of your choice.
9. Click on **NEXT** until the last page of the wizard opens.
 - a. For a conversion to Master/Slave, Stealth and Multi-Master smart architectures, you can configure the *Cloud settings* and *Expert mode*.
 - b. For a conversion to Single-Server smart architecture, you can configure the *Expert mode*.
 - c. For a conversion to Farm smart architecture, you can configure the *Cloud settings*.

For more details regarding the smart architecture configuration, refer to the section [Adding a DNS Smart Architecture](#).

10. Click on **OK** to complete the operation. The report opens and closes. In the column **Type**, the server is now listed as a smart architecture.

Enabling HSM on a DNS Smart Architecture

If you have not enabled the HSM feature upon addition of the smart architecture, you can do it at any time. However, it requires that the smart architecture manages an EfficientIP DNS server, configured as Master, with the feature enabled as well. For more details, refer to the chapter [HSM](#).

To enable HSM on DNS smart architectures

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on **⌵**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
4. Click on **NEXT** until the wizard displays the box **Enable HSM**.
5. Tick the box **Enable HSM** to start signing Master zones managed by the smart architecture using an HSM.
6. Click on **OK** to complete the operation. The report opens and closes. The page **All servers** is visible again. The column **Type** displays your changes.

Handling the Status Locked Synchronization

SOLIDserver provides a consistency check for the smart architectures. Once you configured a smart architecture with the server(s) you want to manage, the smart configuration is checked before it is sent to the physical server(s): this ensures the consistency of the configuration and avoids pushing useless information to the server:

- If the check is conclusive, the information is sent to the server and its status is *Synchronized* in the column *Sync* of the page *All servers*.

- If any error is found, the verification stops and the server **Sync** status changes to **Locked Synchronization** once the page is refreshed. To get a valid synchronization status again, you need to "undo" the latest changes. This action loads a new synchronization and uploads the status accordingly.

Once the server is in *Locked synchronization*, the corrupted configuration file is automatically stored locally on the appliance and available for download in the Local Files Listing. It is named `<server_name>-named.conf`. We advise that you take a look at this file because after the first found error, the check stops and returns the *Locked synchronization* status. So if there are several errors, the status is returned over and over again until the file is conclusive and can be sent to the physical server.

You can check for failure in the configuration file [from the GUI](#) or [via CLI](#).

To check for failure in a DNS configuration file from the GUI

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
3. In the drop-down list **SOLIDserver**, verify that the local appliance is selected. Only the host-name appears with no IP address.
4. In the drop-down list **Services**, select *ipmserver*. The logs appear.
5. In the search engine of the column **Log**, type in *CHECK DNSCONF*. The relevant logs appear, the server name is between brackets.

To check for failure in a DNS configuration file via CLI

1. Open an SSH session.
2. Use the following command to retrieve the list of corrupted files:


```
# ls -la /data1/exports/*-named.conf
```
3. Use the following command to get a precise list of all the errors:


```
# /usr/local/nessy2/bin/named-checkconf /data1/exports/<server_name>-named.conf
```
4. Adjust identified statements, once the check runs again, the *Locked Synchronization* status disappears if you now have a valid configuration.

Deleting a DNS Smart Architecture

At any time, you can decide to stop managing your DNS servers through the smart architectures. You might not need to delete a smart architecture, editing it might be enough. For more details, refer to the section [Changing the Type of a DNS Smart Architecture](#).

Before deleting a smart architecture, keep in mind that:

- Deleting a smart architecture does not delete any data from the physical server: it means that you stop managing the server via the smart architecture. However, **the configuration backup of the smart architecture is deleted**, so if the physical server crashes after the smart architecture deletion, you have to configure everything again manually.
- **You cannot delete a smart architecture if it is still managing DNS servers.**

To delete a DNS smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. If the smart architecture is managing DNS servers, remove them. For more details, refer to the section [Removing a DNS Server from a Smart Architecture](#).
3. Tick the smart architecture you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The smart architecture is no longer listed. All the servers that used to be managed are listed as DNS servers of whatever kind in the list **Type**.

Defining a DNS Smart Architecture as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a smart architecture as one of the resources of a specific group allows the users of that group to manage the architecture in question as long as they have the corresponding rights granted.

Granting access to a smart architecture as a resource also makes every physical server it contains available. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter *Managing Groups*.

Chapter 35. Managing DNS Servers

The server is the highest level of the DNS hierarchy. It allows to resolve host queries and access specific areas of a network, as server can be:

- Authoritative: A server that has authority over a number of domain names and can delegate them.
- Recursive: A server that might contain information, if not it directs the querying host toward the relevant DNS server to solve the query.
- Cache: A server that retrieves information (query results) and keeps it saved in order not to have to query the same information over and over again.

You can create and manage 6 different types of servers: Efficient IP DNS, Efficient IP DNS Package Microsoft DNS (including via AD), Nominum ANS, Generic DNS and Amazon Route 53. You can manage them independently or via a smart architecture. For more details, refer to the chapters [Deploying DNS Smart Architectures](#) and [Managing DNS Smart Architectures](#).

In theory, when a host wishes to access a particular domain, a website for instance, a query is sent to a DNS server that processes the resolution as follows:

1. The DNS client host sends a sequence of queries through a resolver to a recursive DNS server;
2. The recursive server contacts the authoritative servers of the root domain. One of them returns the IP address (an NS record) of the server that has authority over the concerned TLD;
3. The recursive server uses the IP address to connect to the TLD authoritative server and obtain the IP address of the server that has authority over the zone;
4. The recursive server uses the IP address to connect to the zone authoritative server and obtain the queried results;
5. The recursive server sends the results back to the DNS client.

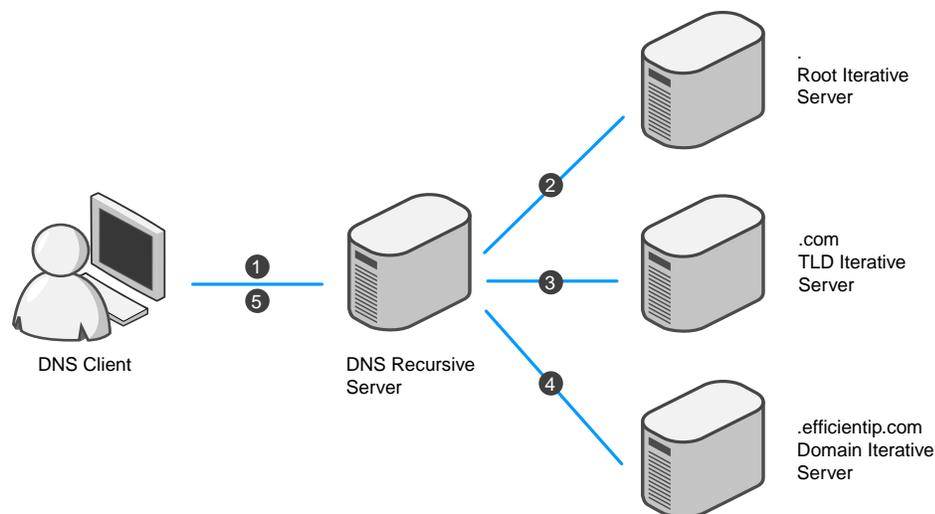


Figure 35.1. Diagram of a DNS query of `www.efficientip.com` via a recursive server

Obviously, such a mechanism would saturate the root zone, which is why a server can combine recursive, cache and or authoritative functionalities.

Browsing DNS Servers

You can manage all your DNS servers from the page *All servers*.

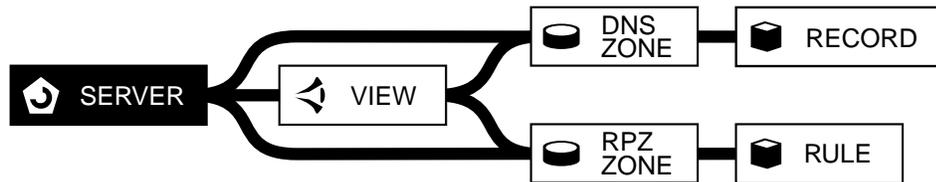


Figure 35.2. The server in the DNS hierarchy

Browsing the DNS Servers Database

To display the list of DNS servers

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. You can filter the list using the column search engines.

To display a DNS server properties page

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **[icon]**. The server properties pages opens.

On the properties page of a physical server, the panel *DNS server statistics* displays all queries statistics in a set of graphs. For more details, refer to the section [Monitoring DNS Servers From their Properties Page](#). Note that if you enabled the service *DNS Guardian*, this page also contains a set of graphs dedicated to DNS Guardian. For more details, refer to the part [Guardian](#).

Customizing the Display on the Page All Servers

Users of the group *admin* can create customized column layouts. The button **[icon] Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Some columns provide more specific information regarding your servers:

Table 35.1. Available columns on the page *All servers*

Column	Description
Guardian	Provides information on the service DNS Guardian on the line of the servers compatible with it. For more details, refer to the part Guardian .
DNSSEC	Indicates if the server is used as a DNSSEC resolver (Yes) or not (No). For more details, refer to the chapter DNSSEC .

Understanding the DNS Server Statuses

The column **Status** provides information regarding the servers you manage.

Table 35.2. DNS server statuses

Status	Description
✔ OK	The server is operational.
ⓘ Timeout	The server does not answer anymore due to a scheduled configuration of the server.
ⓘ License	The license used in SOLIDserver is not compliant with the added server: the license is invalid.
ⓘ Invalid credentials	The SSL credentials are invalid or the server is already managed by another appliance and you need to specify your credentials again. For more details, refer to the section Editing DNS Servers .
ⓘ Syntax error	The server configuration could not be parsed properly.
⊗ Invalid settings	There was a setting error during the server declaration. For instance, some settings were added to a server that does not support them or a smart architecture is not managing any physical server.
ⓘ Invalid time	The server editions performed from the GUI are not pushed to the server because SOLIDserver time and date are incorrect. To correct the time and date refer to the chapter Configuring the Time and Date . In addition, if you are managing Amazon Route 53 servers, you must ensure the appliance time zone is <i>UTC</i> , for more details refer to the section Configuring the User Display Settings .
ⓘ Insufficient privileges	The provided account does not have sufficient privileges to remotely manage the MS server.
⊗ Unmanaged	The server is not available due to a disabling operation.
ⓘ Invalid resolver	SOLIDserver cannot resolve the AWS DNS service. The Amazon services are unreachable and the Amazon Route 53 server cannot be managed. Make sure that the DNS resolvers declared on the page <i>Network configuration</i> are valid.
🕒 Unknown	The server is not synchronized yet.
ⓘ Unknown	An error occurred that SOLIDserver could not identify.
ⓘ Locked synchronization	The server configuration is not viable. This status can be displayed on EfficientIP DNS and EfficientIP DNS Package servers after importing a BIND archive file not properly formed or containing non-supported BIND options. For more details, refer to the section Handling the Status Locked Synchronization .

Note that the column **Sync** changes in accordance with the column **Status**. While the server synchronization is not ✔ OK yet, the column *Sync* might be 🕒 Busy; it can also be in *Locked synchronization*.

The column **Multi-status** provides you with emergency, warning, critical, error or informational messages regarding the server compatibility with Hybrid. For more details, refer to the section [Understanding the Column Multi-Status](#).

Managing EfficientIP DNS Servers

SOLIDserver provides a proprietary DNS server called *EfficientIP DNS* that allows to manage your own server, its configuration and the data it contains.

From the GUI, you can:

- Add EfficientIP DNS servers. For more details, refer to the section [Adding an EfficientIP DNS Server](#).

- Edit EfficientIP DNS servers. For more details, refer to the section [Editing DNS Servers](#).
- Delete EfficientIP DNS servers. For more details, refer to the section [Deleting DNS Servers](#).
- Configure EfficientIP DNS servers. For more details, refer to the chapter [Configuring DNS Servers](#).

Before managing a new server, make sure that the DNS service is correctly started. For more details, refer to the chapter [Configuring the Services](#).

Adding EfficientIP DNS Servers

From the page *All servers*, you can add an EfficientIP DNS Server to manage its configuration, all its data and monitor it. Before adding the server, keep in mind that:

- **The server name is very important.** It is used to publish the NS records of the zone(s) that the server will manage.

Besides, DNS clients must be able to resolve this name when they query the server, so in each zone you must either create an A record or glue records to ensure that they can.

- **Only reachable servers can be added**, each server must be up and running when you add it.
- **The SNMP protocol is no longer supported as managing protocol for a server.** Therefore:
 - You can no longer add a DNS server managed via SNMP.
 - EfficientIP DNS servers prior to version 4.0.x are no longer supported.
 - Your existing servers in version 4.0.x or prior, migrated to 7.0, are still managed via SNMP and listed in the GUI. However, the management of these servers is not detailed in this guide. For more details, refer to the guide *SOLIDserver-Administrator-Guide-5.0.4.pdf*.
- **SSL is used to manage a server while SNMP is used to monitor it.** Therefore:
 - If you manage a DNS server in version 4.0.x legacy, editing it automatically changes the management protocol to use SSL instead of SNMP. This operation is **non-reversible**.
 - You can configure the SNMP monitoring parameters of the server. For more details, refer to the section [Editing the SNMP Monitoring Parameters of an EfficientIP DHCP Server](#).
- When you add a server, a random password is generated to secure the communication between the appliance and the server. The default SSH credentials of the account *admin* are no longer used to manage the server but to generate this random password.

For servers added before the upgrade to version 7.0, switching to this new management system is not automatic. You need to edit the servers. For more details, refer to the section [Editing DNS Servers](#).

- By default, EfficientIP DNS servers embed a hint zone that gathers the IP of all 13 root servers.

To add an EfficientIP DNS server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > EfficientIP DNS**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the following fields to set up the basic server configuration:

Table 35.3. DNS server basic parameters

Field	Description
DNS server name	Type in a DNS resolvable fully qualified domain name (FQDN) for your server. This field is required.
Management IP address	Type in the IPv4 address of your server. This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DHCP. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want. It appears in the column <i>Description</i> of the page <i>All servers</i> . This field is optional.

- If you have modified the SSH password or if the server is already managed by another appliance, tick the box **Configure enrollment parameters**. If not, go to step 7.

Once you have ticked the box, the field "**Admin**" **account password** appears. The default *admin* account password is automatically filled.

- In the field "**Admin**" **account password**, type in your SSH password.
- If you want to edit the server SNMP parameters¹, tick the box **Configure SNMP monitoring parameters**. If not, go to step 8.

Configure the SNMP parameters:

Table 35.4. SNMP parameters used to monitor the server statistics

Field	Description
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
Use TCP	Tick the box if you want to use the TCP protocol instead of the UDP when the network link is not reliable.
SNMP profile	The SNMP profile used to retrieve the statistics. By default, <i>standard v2c</i> is selected. The list contains the default profiles (<i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i>) and the ones you may have created. Each profile has its own level of security and enables the definition of a global security policy. For more details, refer to the section Managing SNMP Profiles .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.

- Tick the box **Enable HSM** to start signing Master zones managed by the server using an HSM. For more details, refer to the chapter [HSM](#).

¹The SNMP protocol parameters are used to monitor and retrieve the server statistics.

- Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 35.5. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The server is listed. The server might appear  *Busy* in the column *Status*. It changes to  *OK* after a while.

During the first DNS server addition, the allow-transfer option is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might need to change the ACL and restrict the option use as it is inherited by the server zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Once added, you can edit your server to secure its data exchanges with SOLIDserver. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).

Editing the SNMP Monitoring Parameters of an EfficientIP DNS Server

Once added to the page *All servers*, you can edit the SNMP monitoring parameters of an EfficientIP DNS server.

The SNMP protocol is no longer supported as managing protocol for a server. If you want to edit the SNMP parameters of a legacy server managed via SNMP, refer to the guide *SOLIDserver-Administrator-Guide-5.0.4.pdf* available on our website ².

To edit the SNMP monitoring parameters of an EfficientIP DNS server

- In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the server of your choice, click on . The properties page opens.
- In the panel **SNMP monitoring properties**, click on . The wizard **SNMP parameters** opens.
- Edit the monitoring parameters according to your needs:

Table 35.6. SNMP parameters used to monitor the server statistics

Field	Description
SNMP version	The version of the SNMP protocol used to retrieve the statistics. It can be either <i>v1</i> , <i>v2c</i> or <i>v3</i> . By default, <i>v2c</i> is selected.
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .

²At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation* and in the folder */docs/5.0.4*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

Field	Description
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.
Use bulk	If you use SNMP version 2 or 3, you can choose to use a bulk transfer of data. This compact SNMP request method accelerates transfers by sending several requests at once. By default, it is set to Yes.
Use TCP	Choose to use the TCP protocol instead of the UDP when the network link is not reliable. By default UDP is used, the drop-down list is set to No.

- Click on **NEXT**. The page **SNMP profile** opens.
- In the drop-down list **SNMP profile**, choose a profile using the same version of the SNMP protocol as the one you selected in the field *SNMP version*.

If you created SNMP profiles, you can choose one of your profiles. They are listed only if they use the same version of the SNMP protocol as the one you selected on the previous page.

Note that the SNMP profiles you can choose from must be configured on the appliance you are currently working with. For more details, refer to the section [Managing SNMP Profiles](#).

- Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

Managing Agentless Microsoft DNS Servers

You can add Microsoft DNS servers to manage them from the page *All servers*. They can be included into an Active Directory (AD) domain or not.

You can reproduce the Microsoft Multi-Master behavior with the smart architecture Multi-Master. This architecture supports Microsoft DNS server, SOLIDserver DNS, BIND server and Nominum's ANS server as well. For more details, refer to the section [Multi-Master Smart Architecture](#).

Once you manage a server, you can also manage its parameters, zones and records.

The management of Microsoft DNS servers is based on Microsoft Remote Procedure Calls (MSRPC) and allows to retrieve and display data in real-time and avoid installing any WinDHCP agent. Microsoft DNS servers with agent are not supported.

Prerequisites

- A Microsoft Window Server 2008, 2008 R2, 2012 R2, 2016 or 2019. The server must:
 - Have the TCP ports 135 and 445 open. They are used by the port mapper interface, the services that indicates to the clients which port provides access to each service.
 - Have Firewall policies that allow traffic between SOLIDserver and the Microsoft servers it manages.
 - In Windows Server 2008, RPC uses by default the dynamic port range 49152-65535. Note that you can reduce the number of available ports, as long as you respect the minimum number of ports required in the range, which is 255, via the netsh tool³.

³For information, refer to <http://support.microsoft.com/kb/929851> .

- The credentials of a member of the groups *DnsAdmins* and *Domain Admins*. Users with insufficient privileges cannot manage the server.
- The service DNS is properly started. For more details, refer to the chapter [Configuring the Services](#).
- The zones of the MS Agentless DNS server must allow the server management IP address in their statement *allow-transfer*.

Limitations

The management of Microsoft DNS servers within SOLIDserver has some limitations. For more details regarding the Microsoft limitations, refer to their documentation.

Server Limitations

- You must refresh manually the DNS server parameters, the list of zones and their parameters. However, the content of the zones is still refreshed automatically every 3600 seconds (by default).
- The AD configuration of the AD integrated DNS servers often includes security settings that prevent the creation or modification of the DNS zones.
- If the parameter *Forward* is set to *!= none* at server level but a list of forwarders is provided anyway, the forwarders are pushed onto the Microsoft DNS server.

ACL Limitations

- Microsoft processes as follows the ACL allow-update:

Table 35.7. Limitations for the ACL allow-update

Allow-update set to	Microsoft behavior
<i>admin;any;</i>	The update rights are granted to anyone.
<i>any;</i>	The update rights are granted to anyone.
<i>admin;</i>	If the zone is AD integrated, the update is changed to <i>Secure Only</i> .
Any other parameter	The update is impossible, the allow-update is set to <i>no update</i> .

- Microsoft processes as follows the ACL allow-transfer:

Table 35.8. Limitations for the ACL allow-transfer

Allow-transfer set to	Microsoft behavior
<i>any;</i>	The transfer rights are granted to anyone.
No parameter is set	The transfer rights are granted to anyone.
<i>none;</i>	The transfer rights are not granted to anyone.
<i>eip_ns_only;</i>	Only the transfer of the zone NS resource records is granted.
Any other parameter	If ACLs are set, they are ignored.
Any other parameter	If IP addresses are listed, the allow-transfer is granted to the <i>Specified IP Address List</i> and to the zone NS resource records.

Zone Limitations

- The zones *e164.arpa* and *ip6.int* (deprecated reverse mapping name space) are not supported by Microsoft.
- You cannot create *forward* zones with the forwarding parameter set to *None* on Microsoft servers.

- If nothing is specified during the Notify configuration then by default, the notify is set to *NS only*.
- You cannot edit the *AD replication* behavior set on the zones of an AD integrated Microsoft server. Once set, you cannot edit or remove the replication behavior selected, you must delete the zone and recreate it with the configuration that suits your needs.

Resource Record Limitations

- Microsoft servers only support the records A, AAAA, CNAME, MX, NS, PTR, SOA, SRV and TXT.

Adding Agentless Microsoft DNS Servers

Before adding an Agentless Microsoft DNS server make sure you meet the [Prerequisites](#).

If your Microsoft DNS server is integrated to an AD with several forests, you can use the *Expert mode* during the server addition to specify AD domain you want to authenticate.

To add an agentless Microsoft DNS server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > Microsoft DNS (agentless)**. The wizard **Add a DNS server** opens.
3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Configure the server basic parameters:

Table 35.9. DNS server basic parameters

Field	Description
DNS server name	Type in an FQDN name for the server. This field is required.
Management IP address	Type in the IPv4 address of the Microsoft DNS server you want to manage ^a . This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DHCP. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want. This field is optional.

^aWith the proper [Configuring the Network](#), if you enter the name of your DNS server in this field and click on **SEARCH**, the IP address is retrieved from the DNS and displayed.

5. Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 35.10. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .

Field	Description
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

6. Click on . The last page of the wizard opens.
7. In the field **User**, type in the name of a user with sufficient management privileges over the Windows DNS server.
8. In the field **Password**, type in the corresponding password.
9. If your Windows DNS server is integrated to Active Directory and contains several forests:
 - a. Tick the box **Expert mode (AD)**. The field *AD domain* appears
 - b. In the field **AD domain**, type in the full domain or subdomain.
10. Click on to complete the operation. The report opens and closes. The server is listed, the column **Version** indicates the Microsoft server version.

Managing BIND DNS Servers

SOLIDserver provides its software versions through native packages of operating system. Installing the DNS package allows you to use the DNS module of SOLIDserver at the best of its potential on Linux: it allows you to manage your BIND server through an EfficientIP DNS server, which incidentally provides all the options that come with it (statistics retrieved via SNMP...).

The addition of Linux Packages v4 that respected the SNMP protocol and worked with SSL is no longer supported. If you migrated your database with such servers, refer to the guide *SOLIDserver-Administrator-Guide-5.0.4.pdf* available on our website ⁴.

Managing EfficientIP BIND Linux Packages

In the sections below are a set of procedures to successfully install the DNS packages, formerly known as EfficientIP BIND Linux Packages v5, on Linux Debian and CentOS/RedHat.

Installing the EfficientIP DNS Package for Linux Debian 9 and prior - 64 bits

You must take into account the [prerequisites](#) before [installing](#) a DNS Debian package.

Prerequisites

- The DNS package file, *ipmdns-y.y.debianxx-amd64.deb*, whose name provides you with a number of information separated by hyphens: the type of package (*ipmdns*, so a DNS package), the version of SOLIDserver (*y.y.y*); the version of Debian (*debianxx* where *xx* is *x dot x*) and finally the Debian architecture (*amd64*).

In the procedure below, this file is referred to as `<ipmdns-y.y.y-debianxx-amd64.deb>`.

⁴At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation* and in the folder */docs/5.0.4*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

- The EfficientIP BIND package platform must have at least 20 Mb of free disk space.
- The EfficientIP BIND package may need certain libraries of your operating system, you must have a *shell* access with *root* login in local, via *ssh*⁵ on the server to be installed.
- You must make sure that no other DNS/DHCP service on your Linux is running : it would interfere with the BIND/ISC package installation.
- You must make sure that SOLIDserver and Debian are set to the same time and date.
- You must make sure that Apache server is up-to-date.
- You must make sure that the service *dbus* is installed.
- You must make sure that HTTPS (port 443) and the DNS service (port 53) are not blocked by a network filtering process (firewall).

If your Apache configuration already uses the port 443, you have to create an additional IP-based VirtualHost dedicated to the DNS/DHCP management.

Installing the EfficientIP DNS Package

You can install the EfficientIP DNS Package on Debian Linux.

If you have not installed the DHCP packages yet, you need to:

1. Follow the procedure [To install the EfficientIP DNS Package on Debian](#).
2. Follow the procedure [To complete the DNS package installation on Debian if the DHCP package is not installed](#).

If you already installed the DHCP packages, you only need to follow the procedure [To install the EfficientIP DNS Package on Debian](#) below.

The procedure below includes the commands that make the web services configurable.

To install the EfficientIP DNS Package on Debian

1. Open an SSH session.
2. Stop and disable your DNS software: If you are using NSD or Unbound, refer to the related proprietary documentation. In the case of BIND, use the following commands:

```
# service bind9 stop
# update-rc.d -f bind9 remove
```

3. Install the dependency packages, **ONLY** if you have not installed the EfficientIP DHCP package, using the following commands:

- a. For Debian 7 and prior:

```
# apt-get install php5
# apt-get install sudo
# apt-get install snmpd
```

- b. For Debian 8 and Debian 9:

```
# apt-get install php
# apt-get install sudo
# apt-get install snmpd
```

⁵You could also connect via *telnet* but, for security purposes, we recommend that you favor *ssh*.

4. Install the EfficientIP DNS package, using the following command:

```
# dpkg -i <ipmdns-y.y.y-debianxx-amd64.deb>
```

5. Make the web services configurable: in the directory `/etc/sudoers.d`, create the file `ipmdns` containing the line below.

```
www-data ALL = NOPASSWD: /usr/local/nessy2/script/install_named_conf.sh, \  
/usr/local/nessy2/script/push_default_zone_params.sh, \  
/usr/local/nessy2/script/push_dnssec_keys_zones.sh, \  
/usr/local/nessy2/script/move_dnszone_file.sh, \  
/usr/local/nessy2/script/restore_named_conf.sh, \  
/usr/local/nessy2/script/delete_zone_file.sh, \  
/usr/local/nessy2/script/restore_zone_file.sh, \  
/usr/local/nessy2/script/install_keytab.sh, \  
/usr/local/nessy2/bin/rndc
```

6. Set the users access rights as follows:

```
# chmod 440 /etc/sudoers.d/ipmdns
```

Note that you can change the password `admin` of the web service using the command below:

```
# htpasswd -c /usr/local/nessy2/www/php/cmd/dns/.htpasswd admin
```

If you have not installed the DHCP package or are not planning on installing it, you must now follow the procedure below.

To complete the DNS package installation on Debian if the DHCP package is not installed

1. If relevant, open an SSH session.
2. Allow SNMP access to the DNS statistics: append the file `/etc/snmp/snmpd.conf` with the following line.

```
view systemonly included .1.3.6.1.4.1.2440
```

3. Start the snmp daemon, using the following command:

```
# service snmpd start
```

4. Create a self-signed certificate for Apache, using the following commands:

```
# cd /etc/apache2  
# openssl genrsa -des3 -out server.key 4096  
# openssl req -new -key server.key -out server.csr  
# openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt  
# openssl rsa -in server.key -out server.key.insecure  
# mv server.key server.key.secure  
# mv server.key.insecure server.key
```

5. Activate the SSL mode in Apache using the following command:

```
# a2enmod ssl
```

6. Make sure that a symbolic link to the default VirtualHost SSL configuration file is located in the folder `sites-enabled/`. If not, use the following command.

- a. For Debian 7 and prior:

```
# ln -s /etc/apache2/sites-available/default-ssl /etc/apache2/sites-enabled/
```

- b. For Debian 8 and Debian 9:

```
# a2ensite default-ssl
```

7. Configure the web services.

- a. For Debian 7 and prior, in the file `/etc/apache2/sites-enabled/default-ssl`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

```
<VirtualHost *:443>

ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    AllowOverride All
</Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/apache2/server.crt
SSLCertificateKeyFile /etc/apache2/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

- b. For Debian 8 and Debian 9, in the file `/etc/apache2/sites-enabled/default-ssl.conf`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

```
<VirtualHost *:443>

ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
    <Directory /usr/local/nessy2/www/php>
        Require all granted
        AllowOverride Authconfig
        Options Indexes FollowSymLinks
    </Directory>

# Please note that the following, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/apache2/server.crt
SSLCertificateKeyFile /etc/apache2/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM
```

```
# Please note that the following, from "SetEnvIf" to "force-response-1.0", represents and
# must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0
</VirtualHost>
```

8. Disable the default site in Debian Apache configuration using the following commands.

a. For Debian 7 and prior:

```
# cd /etc/apache2/sites-enabled
# unlink 000-default
```

b. For Debian 8 and Debian 9:

```
# a2dissite 000-default
```

9. Restart Apache using the following command:

```
# service apache2 restart
```

10. Make sure that the *ipmdns* package is running using the following command:

```
# service ipmdns status
```

If it is not running, use the following command:

```
# service ipmdns start
```

Once the configuration is complete, you can add an EfficientIP Package DNS server to manage your BIND server from SOLIDserver GUI. Refer to the procedure in the section [Adding a BIND DNS Server](#) for more details.

Installing the EfficientIP DNS Package for Linux CentOS/RedHat 7 and prior - 64 bits

You must take into account the [prerequisites](#) before [installing](#) a DNS CentOS/RedHat package.

Prerequisites

- The DNS package file, *ipmdns-y.y.y-redhatx.x86_64.rpm*, whose name provides you with a number of information separated by hyphens or a point: the type of package (*ipmdns*, so a DNS package), the version of SOLIDserver (*y.y.y*); the version of RedHat (*redhatx*) and finally the RedHat architecture (*x86_64*).

In the procedure below, this file is referred to as `<ipmdns-y.y.y-redhatx.x86_64.rpm>`.

- The EfficientIP BIND package platform must have at least 20 Mb of free disk space.
- The EfficientIP BIND package may need certain libraries of your operating system, you must have a *shell* access with *root* login in local, via *ssh*⁶ on the server to be installed.
- You must make sure that no other DNS/DHCP service on your Linux is running : it would interfere with the BIND/ISC package installation.
- You must make sure that SOLIDserver and RedHat/CentOS are set to the same time and date.
- You must make sure that Apache server is up-to-date.

⁶You could also connect via *telnet* but, for security purposes, we recommend that you favor *ssh*.

- You must make sure that HTTPS (port 443) and the DNS service (port 53) are not blocked by a network filtering process (firewall).

If your Apache configuration already uses the port 443, you have to create an additional IP-based VirtualHost dedicated to the DNS/DHCP management.

Installing the EfficientIP DNS Package

You can install the EfficientIP DNS Package on both RedHat and CentOS Linux.

If you have not installed the DHCP packages yet, you need to:

1. Follow the procedure [To install the EfficientIP DNS Package on RedHat and CentOS](#).
2. Follow the procedure [To complete the DNS package installation on RedHat/CentOS if the DHCP package is not installed](#).

If you already installed the DHCP packages, you only need to follow the procedure [To install the EfficientIP DNS Package on RedHat and CentOS](#) below.

The installation procedure below also includes the commands that make the web services configurable.

To install the EfficientIP DNS Package on RedHat and CentOS

1. Open an SSH session.
2. Stop and disable your DNS software, using the following commands:

```
# service named stop
# chkconfig named off
```

3. Install the dependency packages, ONLY if you have not installed the EfficientIP DHCP package, using the following commands:

```
# yum install mod_ssl php php-pdo sudo net-snmp
```

4. For RedHat/CentOS version 5 or 6, you must install the dependency packages:

- a. Install the repository *EPEL*:

```
# yum install epel-release
```

- b. Install the library *GeoIP* in the repository *epel-release*:

```
# yum install GeoIP
```

5. Install EfficientIP DNS package, using the following command:

```
# rpm -ivh <ipmdns-y.y.y-redhatx.x86_64.rpm>
```

6. Make the web services configurable: in the directory */etc/sudoers.d*, create the file *ipmdns* containing the line below.

```
apache ALL = NOPASSWD: /usr/local/nessy2/script/install_named_conf.sh, \
/usr/local/nessy2/script/push_default_zone_params.sh, \
/usr/local/nessy2/script/push_dnssec_keys_zones.sh, \
/usr/local/nessy2/script/move_dnszone_file.sh, \
/usr/local/nessy2/script/restore_named_conf.sh, \
/usr/local/nessy2/script/delete_zone_file.sh, \
/usr/local/nessy2/script/restore_zone_file.sh, \
```

```
/usr/local/nessy2/script/install_keytab.sh, \  
/usr/local/nessy2/bin/rndc
```

7. Set the users access rights as follows:

```
# chmod 440 /etc/sudoers.d/ipmdns
```

Note that you can change the password *admin* of the web service using the command below:

```
# htpasswd -c /usr/local/nessy2/www/php/cmd/dns/.htpasswd admin
```

If you have not installed the DHCP package or are not planning on installing it, you must now follow the procedure below.

To complete the DNS package installation on RedHat/CentOS if the DHCP package is not installed

1. If relevant, open an SSH session.
2. Disable the firewall using the following commands.

- a. For RedHat/CentOS 6 and prior:

```
# service iptables stop  
# chkconfig iptables off
```

- b. For RedHat/CentOS 7:

```
# service firewalld stop  
# chkconfig firewalld off
```

3. Disable selinux. In the file */etc/selinux/config*, modify the line *SELINUX=enforcing* to match the following one:

```
SELINUX=disabled
```

Note that changing the selinux policy requires you to restart the system.

4. Reboot the system to take into account the selinux policy changes :

```
# reboot
```

5. In the file */etc/sudoers*, disable *requiretty* by making it a comment as follows:

```
#Defaults    requiretty
```

6. Allow SNMP access to the DNS statistics. In the file */etc/snmp/snmpd.conf*, in the section entitled Access Control, enter the lines:

```
master agentx  
view systemview included .1.3.6.1.4.1.2440  
#You may need to specify another view, AllView or a custom one,  
#if you edited the default SNMP configuration.
```

7. Start the SNMP daemon, using the following command:

```
# service snmpd start
```

8. Create a self-signed certificate for Apache, using the following commands:

```
# cd /etc/httpd  
# openssl genrsa -des3 -out server.key 4096  
# openssl req -new -key server.key -out server.csr  
# openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt
```

```
# openssl rsa -in server.key -out server.key.insecure
# mv server.key server.key.secure
# mv server.key.insecure server.key
```

9. Configure the web services. In the file `/etc/httpd/conf.d/ssl.conf`, replace the FULL section `<VirtualHost *:443>` with the configuration below.

- a. For RedHat/CentOS 6 and prior:

```
<VirtualHost *:443>
ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    AllowOverride All
</Directory>

# Please note that the following data, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/httpd/server.crt
SSLCertificateKeyFile /etc/httpd/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following data, from "SetEnvIf" to "force-response-1.0", represents
# and must be written on a single line.
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0

</VirtualHost>
```

- b. For RedHat/CentOS 7:

```
<VirtualHost *:443>
ServerName 127.0.0.1
DocumentRoot /usr/local/nessy2/www/php
<Directory /usr/local/nessy2/www/php>
    Require all granted
    AllowOverride Authnconfig
    Options Indexes FollowSymLinks
</Directory>

# Please note that the following data, from "php_admin_value" to "site:/usr/local/share/pear",
# represents and must be written on a single line.
php_admin_value include_path
/usr/local/nessy2/www/php/include:/usr/local/nessy2/lib/php:/usr/local/nessy2/www/site:/usr/local/share/pear
php_admin_value file_uploads 1
php_admin_value upload_max_filesize 300000000
php_admin_value post_max_size 300000000
php_admin_value memory_limit 150000000
php_admin_value register_globals 0
php_admin_value short_open_tag 1
php_admin_value safe_mode 0
php_admin_value magic_quotes_gpc 0

SSLEngine on
SSLCertificateFile /etc/httpd/server.crt
SSLCertificateKeyFile /etc/httpd/server.key
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM

# Please note that the following data, from "SetEnvIf" to "force-response-1.0", represents
# and must be written on a single line.
```

```
SetEnvIf User-Agent ".*MSIE.*" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0
</VirtualHost>
```

- Restart Apache using the following command:

```
# service httpd restart
```

- Make sure that the *ipmdns* package is running using the following command:

```
# service ipmdns status
```

If it is not running, use the following command:

```
# service ipmdns start
```

Once the configuration is complete, you can add an EfficientIP Package DNS server to manage your BIND server from SOLIDserver GUI. Refer to the procedure in the section [Adding a BIND DNS Server](#) for more details.

Upgrading Packages

No matter the package version, to upgrade packages:

- You must **uninstall your current packages**.
- You must **install the new package** following the prerequisites and procedures detailed above in the section [Managing EfficientIP BIND Linux Packages](#).

Adding BIND DNS Servers

Once you successfully installed your EfficientIP BIND Linux package, you need to install an EfficientIP DNS package and configure it according to your needs to manage BIND servers through the GUI.

Note that only reachable servers can be added, the server must be up and running when you add it.

To add a BIND DNS server for a Linux package

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- In the menu, select **+ Add > Server > EfficientIP DNS Package**. The wizard **Add a DNS server** opens.
- If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the following fields to set up the basic server configuration:

Table 35.11. DNS server basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.

Field	Description
Management IP address	Type in the IPv4 address of your server. This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DHCP. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page All servers. This field is optional.

- If you modified the SSH login and password, tick the box **Configure SSL parameters**. If not, go to step 6.

Once you ticked the box, the fields **Login** and **Password** appear. By default they both contain *admin*, edit them to make sure that the SSL credentials match your SSH credentials.

- If you want to edit the server SNMP parameters⁷, tick the box **Configure SNMP monitoring parameters**. If not, go to step 7.

Once you ticked the box, the following fields appear:

Table 35.12. SNMP parameters used to monitor the server statistics

Field	Description
SNMP port	The port used to retrieve the server statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
Use TCP	Tick the box if you want to use the TCP protocol instead of the UDP when the network link is not reliable.
SNMP profile	The SNMP profile used to retrieve the statistics. By default, <i>standard v2c</i> is selected. The list contains the default profiles (<i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i>) and the ones you may have created. Each profile has its own level of security and enables the definition of a global security policy. For more details, refer to the section Managing SNMP Profiles .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5s.

- Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 35.13. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .

⁷The SNMP protocol parameters are used to monitor and retrieve the server statistics.

Field	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **OK** to complete the operation. The report opens and closes. The list is visible again. The server appears in the list with status  *Busy*. It changes to  *OK* after a while.

During the first DNS server addition, the allow-transfer option is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might need to change the ACL and restrict the option use as it is inherited by the server zones. For more details, refer to the chapter [Limiting Zone Transfers at Server Level](#).

Once the EfficientIP Package server is added, you can manage your BIND server in Linux from the GUI.

Managing Generic DNS Servers

SOLIDserver can manage generic DNS servers that are not: EfficientIP, Microsoft, or EfficientIP BIND packages. However, the possibilities for managing such servers are more restricted than they are for other DNS servers. If these DNS servers support dynamic DDNS updating as described in RFC 2136, the contents of their zones can be administered by SOLIDserver management console. The Generic DNS management imports the data through zone transfers from the remote DNS server. The remote DNS server must allow zone transfers. This can be performed by configuring the IP address of the SOLIDserver management to which the data should be imported.

Before managing a new server make sure that the DNS service is correctly started. For more details, refer to the chapter [Configuring the Services](#).

Adding Generic DNS Servers

To fully configure and manage a Generic DNS server, you need to follow three procedures in order to:

- Add a *Generic* server to the list All servers.
- Configure its TSIG parameters if need be.
- Add the DNS zones that should be managed through the server.

To add a generic DNS server

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- In the menu, select **+ Add > Server > Generic DNS**. The wizard **Add a DNS server** opens.
- Fill in the fields below:

Table 35.14. DNS server basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Management IP address	Type in the IP address of your server. This field is required.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page All servers. This field is optional.

- Click on **OK** to complete the operation. The server is listed in the page **All servers**.

Once added, you can edit your server to secure its data exchanges with SOLIDserver. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).

To add generic DNS server zones

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- Click on the name of the Generic server. The page **All zones** opens.
- In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
- If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Click on **NEXT**. The next page opens.
- In the list **DNS zone type**, select *Master*.
- In the list **DNS zone resolution**, select *Name*.
- Click on **NEXT**, the next page opens.
- Fill in the fields according to the table below:

Table 35.15. Master zone creation fields

Field	Description
Name	Type in the zone name you chose. It should strictly conform with the syntax given in RFC1034.
Space	Select the space tied to that zone. Assigning IP addresses in the selected space updates the DNS zone you are creating.

- Click on **NEXT**. The last page of the wizard opens.
- The fields on that page are automatically filled. However you can edit them following the table below. All the fields are required.

Table 35.16. Zone advanced parameters

Field	Description
Primary server	Specify the primary Master server for the zone.
Responsible	Specify the administrator email address for the zone.
Serial number	The zone serial number. It is automatically incremented for each zone change.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expire	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

You can set the value by default for the parameters above, except for the *Primary server* and *Serial number*. For more details, refer to the procedure [To configure the default SOA parameters of Master zones](#) below.

- Click on to complete the operation. The report opens and closes. The report opens and closes. The zone is listed and marked *Delayed create* before being marked **OK**.

Managing Nominum ANS Servers

In addition to traditional DNS servers, SOLIDserver allows you to manage Nominum authoritative name servers (ANS).

To fully configure and manage a Nominum DNS server, you first need to prepare the ANS security key, or password, related to said server and follow the procedures below in order to:

- Add a Nominum server to the list All servers.
- Add the DNS zones that should be managed through the server.

Before managing a new server make sure that the DNS service is correctly started. For more details, refer to the chapter [Configuring the Services](#).

Adding Nominum ANS Servers

From the DNS All servers list you can add a Nominum ANS server.

To add a Nominum ANS server

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- In the menu, select **Add > Server > Nominum ANS**. The wizard **Add a DNS server** opens.
- If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on . The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Fill in the fields below:

Table 35.17. DNS server basic parameters

Field	Description
DNS server name	Type in a FQDN name for your server. This field is required.
Management IP address	Type in the IP address of your server. This field is required.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page All servers. This field is optional.
ANS key	Type in the security key, or password, configured on the Nominum server. This field is required.

- Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 35.18. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The server is listed on the page **All servers**.

Once added, you can edit your server to secure its data exchanges with SOLIDserver. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).

Adding Zones to a Nominum ANS Server

Once you created a Nominum ANS server, you can add your zones from its All zones list.

To add ANS server zones

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- Click on the name of the Nominum ANS server. The page **All zones** opens.
- In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
- If you or your administrator created classes at zone level, in the list **DNS zone class** select a class or *None*. Click on . The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the list **DNS zone type**, select *Master*.
- In the list **DNS zone resolution**, select *Name*.
- Click on . The next page of the wizard opens.
- Fill in the fields according to the table below:

Table 35.19. Master zone creation fields

Field	Description
Name	Type in the zone name you chose. It should strictly conform with the syntax given in RFC1034.
Space	Select the space tied to that zone. Assigning IP addresses in the selected space updates the DNS zone you are creating.

- Click on . The last page of the wizard opens.
- The fields on that page are automatically filled. However you can edit them following the table below. All the fields are required.

Table 35.20. Zone advanced parameters

Field	Description
Primary server	Specify the primary Master server for the zone.
Responsible	Specify the administrator email address for the zone.
Serial number	The zone serial number. It is automatically incremented for each zone change.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expire	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

You can set the value by default for the parameters above, except for the *Primary server* and *Serial number*. For more details, refer to the procedure [To configure the default SOA parameters of Master zones](#) below.

- Click on to complete the operation. The report opens and closes. The zone is listed and is marked *Delayed create* before being marked **OK**.

Managing Amazon Route 53 Servers

SOLIDserver supports the management of Route 53 DNS servers. These servers provide a cloud based Anycast DNS service available for Amazon Web Service (AWS) account owners.

From the DNS All servers page, you can add your Route 53 server using your AWS account credentials. Once listed in the GUI, you can add, edit and/or delete the server zones and resource records.

Before managing a new server make sure that the DNS service is correctly started. For more details, refer to the chapter [Configuring the Services](#).

Prerequisites

- You must have an AWS account with a subscription to the service Amazon Route 53.
- You must have an Amazon user that should only have access to the service Amazon Route 53 and sufficient rights to manage the service.

For security reasons, we strongly recommend that you create this user through the Amazon IAM module. Besides, the user needs an Access Key ID and a Secret Access Key to manage the service, you should also generate these keys using the module IAM.

- Before adding the server to SOLIDserver, you should have your AWS account Access Key ID and Secret Access Key ready.
- You must use the UTC system to ensure that SOLIDserver and the AWS account are set at the same time. For more details, refer to the procedure [To configure the user settings](#).
- Make sure SOLIDserver is able to contact the AWS services using REST and HTTPS. Otherwise, even if you add the server to the GUI, it stays in *Timeout*.

Limitations

Amazon Route 53 Server Limitations

- SOLIDserver only supports unique zone names within one server.

If your Amazon Route 53 server contains several zones named the same way, none of them can be managed from the GUI.

- DNSSEC is not supported on Amazon Route 53 servers.
- The edition options are limited:
 - You can only edit the Main properties and Group access panels of an Amazon Route 53 server managed outside a smart architecture.
 - Once an Amazon Route 53 server is managed via a smart architecture, you can only edit the panel *Group access* on the server properties page.
 - Converting master zones to slave on Amazon Route 53 servers deletes all the records of the original master zone.
- Some options are not supported on Amazon Route 53 servers:
 - You cannot edit the access control of an Amazon Route 53 server or its zone.
 - You cannot edit the options of an Amazon Route 53 server or its zone.

These options cannot be edited either from SOLIDserver or from the service Amazon Route 53 itself.

- Some AWS options are not supported by SOLIDserver: the record options *Health Check* and *Routing Policy* are not supported.

If any AWS zone contains records configured with these options, the entire zone is not synchronized.

Amazon Route 53 Zone and Record Limitations

Before managing an AWS server from the GUI, we recommend that you check the prerequisites, specificities and limitations of the AWS zones and records you are managing. They are all listed in the sections [Managing Amazon Route 53 Zones](#).

Adding Amazon Route 53 Servers

When adding an Amazon Route 53 server, you do not need to provide an IP address. Thanks to the AWS account credentials, SOLIDserver identifies your Amazon Route 53 server and communicates directly with it using REST protocol.

If you have several AWS accounts, you can manage all your Amazon Route 53 servers from the GUI.

Once you comply with the [Prerequisites](#), you can follow the procedure below.

To add an Amazon Route 53 server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > Amazon Route 53**. The wizard **Add a DNS server** opens.

3. If you or your administrator created classes at the server level, in the list **DNS server class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. Fill in the following fields to set up the basic server configuration:

Table 35.21. DNS server basic parameters

Field	Description
DNS server name	Specify the FQDN name of your server. This field is required.
Isolated	Tick this box if you want to prevent the server configuration from updating the module IPAM or DHCP. Any advanced property set via the drop-down list <i>Advanced properties</i> has to be applied to the server later on. This option is mainly useful when dealing with migrations. Keep in mind that the server still receives data if your network configuration allows it. Before unticking the box, make sure that the configuration you set suits your needs.
Description	Type in a description if you want, it appears in the column <i>Description</i> of the page <i>All servers</i> . This field is optional.

5. Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 35.22. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

6. Click on **NEXT**. The last page of the wizard opens.
7. In the field **Access Key ID**, type in your AWS account Access Key ID.
8. In the field **Secret Access Key**, type in the corresponding Secret Access Key.
9. Click on **OK** to complete the operation. The server is listed on the page **All servers**.

Once you added an Amazon Route 53 server:

- The server stays in *Invalid time* if SOLIDserver and the AWS account are not set at the same time. SOLIDserver must be set to UTC. For more details, refer to the procedure [To configure the user settings](#).
- If you edit the content of an Amazon Route 53 server directly from the AWS account, you need to synchronize the server in SOLIDserver GUI. For more details, refer to the procedure [To synchronize DNS servers](#).
- If your internal network policies prevent you from accessing Amazon services directly, you might need to set a proxy server via a registry database entry to handle Amazon Route 53 requests, as described below.

To set a proxy server for Amazon Route 53

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search field **Name**, type in *module.dns.route53.proxy* and filter the list.
4. In the column **Value**, click on the corresponding value. The wizard **Registry database Add an item** opens.
5. In the field **Value**, type in the server configuration as follows: `<hostname_or_ip>:<port>`. If need be, you can also define the access credentials as follows: `<username>:<password>@<hostname_or_ip>:<port>`.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the key is listed.

Managing Amazon Route 53 Servers With a Smart Architecture

All the smart architectures can manage Amazon Route 53 servers, either alone or in addition to other physical servers:

- **Single-server:** no requirement.
- **Multi-Master:** in addition to at least one Master server and with no Slave server.
- **Master/Slave:** in addition to at least one Master server and one Slave server.
- **Farm:** in addition to at least one Master server and one Slave server.
- **Stealth:** in addition to at least one Hidden-master server and one Pseudo-master server.

Before adding your server to a smart:

- We recommend that you tick the AWS server and generate the report **Route 53 Incompatibilities**. That way you can edit the server if need be and make sure that the server you manage through the smart is properly configured. For more details, refer to the section [DNS Server Reports](#).
- Make sure your AWS server status is  OK.

Once you manage an Amazon Route 53 server with a smart architecture:

- It has the role *Cloud*.
- It cannot be set as *Master* or *Slave*.
- Only the Master zones of the other physical servers managed through the smart architecture are replicated on the Amazon Route 53 server.
- Only the supported configuration settings of the other physical servers managed through the smart architecture are pushed to the Amazon Route 53 server. For more details, refer to the section [Limitations](#).
- The smart architecture management remains the same: the smart architecture replicates all the server options and content from one server to the other. However, all the configuration settings that are incompatible with Amazon Route 53 servers, or not supported, are not replicated and displayed in the *Multi-Status* column.

- Any changes performed from the GUI - zones or records addition, edition or deletion - automatically refresh the server, it then replicates the information to your AWS account server.
- You cannot delete the *awsdns* NS records from the list All RRs of the smart architecture. The resolution of the Amazon Route 53 server zones rely on them.

Adding an Amazon Route 53 Server When Creating a Smart Architecture

You can add an Amazon Route 53 server to the management of the smart while creating the smart architecture.

This allows to automate the synchronization of the AWS server content. Therefore, all your zones and records are retrieved by the smart and replicated to all the physical servers you might manage as well with the new smart.

To add an Amazon Route 53 server to a smart architecture you are creating

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the menu, select **+ Add > Server > DNS smart architecture**. The wizard **Add a DNS server** opens.
3. Fill in the basic parameters. Click on **NEXT**. The next page of the wizard opens. For more details, refer to the section [Adding a DNS Smart Architecture](#).
4. In the list **DNS smart architecture**, select an architecture. Click on **NEXT**. The page **DNS servers role configuration** opens.
5. If you are creating a *Single-Server* smart architecture:
 - a. In the drop-down list **Available DNS servers**, the AWS server is listed.
 - b. Click on **+ MASTER** to add it to the **Master DNS server(s) list** and manage it via the smart architecture. Being an AWS server, it actually is not Master, it is the only server of the smart. For more details regarding the box **Expert mode**, refer to the section [Single-Server Smart Architecture](#).
6. If you are creating a *Master/Slave*, *Stealth* or *Multi-Master* smart architecture:
 - a. At the bottom of the page, tick the box **Cloud settings**. Two fields appear.
 - b. In the drop-down list **Available Cloud DNS servers**, select the Amazon Route 53 server of your choice.
 - c. Click on **+ CLOUD**. The selected server is moved to the **Cloud DNS servers list**.
 - d. If you want to add other servers to the smart architecture, refer to the section [Adding a DNS Smart Architecture](#).
7. If you are creating a *Farm* smart architecture:
 - a. At the bottom of the page, tick the box **Cloud settings**. Two fields appear.
 - b. In the drop-down list **Available Cloud DNS servers**, select the Amazon Route 53 server of your choice.
 - c. Click on **+ CLOUD**. The selected server is moved to the **Cloud DNS servers list**.
 - d. If you want to add other servers to the smart architecture, refer to the section [Adding a DNS Smart Architecture](#).

- e. Click on **NEXT**. The page **Advanced settings** opens, finish the configuration following the procedure in the section [Farm Smart Architecture](#).
8. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is listed, if you do not see the servers it manages, click on .

Adding an Amazon Route 53 Server in an Existing Smart Architecture

If you add an Amazon Route 53 server to an existing smart architecture managing other servers, keep in mind that the zones it contains rely on a specific set of NS records to run properly. Without all these records, the zones are not viable.

Therefore, the smart architecture automatically:

- Retrieves the content of the AWS server.
- Replicates the content of your AWS server (all the additional NS records) on all the physical servers managed. The All RRs list of all the servers should therefore contain *awsdns* records.

To add an Amazon Route 53 server into an existing smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
4. If need be, edit the server basic parameters. For more details, refer to the section [Adding a DNS Smart Architecture](#). Click on **NEXT**. The next page of the wizard opens.
5. In the list **DNS smart architecture**, edit the type of DNS smart architecture if need be. Click on **NEXT**. The page **DNS servers role configuration** opens.
6. At the bottom of the page, tick the box **Cloud settings**. Two fields appear.
7. In the drop-down list **Available Cloud DNS servers**, select the Amazon Route 53 server of your choice. Only the servers that are not already managed by a smart architecture are listed.
8. Click on **+ CLOUD**. The selected server is moved to the **Cloud DNS servers list**.
9. If you are editing a Farm architecture or if you configured NS records on another architecture, click on **NEXT**. The page **Advanced settings** opens. For more details, refer to the relevant smart architecture addition procedure in the section [Adding a DNS Smart Architecture](#).
10. Click on **OK** to complete the operation. The report opens and closes. The synchronization, sends the physical servers information to your AWS account server. You can display the smart architecture physical servers on the list **All servers** using the  button in the upper right corner.

You can check that the replication is properly performed on the page All RRs of the smart architecture. They should be all listed and  **OK**.

On the page All zones, the column **Multi-status** indicates any replication problems.

Removing an Amazon Route 53 Server from a Smart Architecture

At any time, you can stop managing an AWS server via a smart architecture.

To successfully remove an AWS server from a smart:

1. Follow the procedure [To remove an Amazon Route 53 server from a smart architecture](#).
2. If your smart architecture manages other servers, you should remove the *awsdns* NS records as they are no longer relevant for the smart architecture, as detailed in the procedure [To delete Amazon Route 53 NS resource records](#).

To remove an Amazon Route 53 server from a smart architecture

1. In the sidebar, go to  **DNS** > **Servers**. The page **All servers** opens.
2. At the end of the line of the smart architecture of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on . The wizard **Edit a DNS server** opens.
4. Click on . The smart architecture dedicated page opens.
5. Click on . The page **DNS servers role configuration** opens.
6. Depending on the smart architecture you are editing, in the field **Master DNS server(s) list** or **Cloud DNS servers list**, select the server to remove and click on .

The server is moved back to the list **Available DNS servers** or **Available Cloud DNS servers**.

7. If you are editing a Farm architecture or if you configured NS records on another architecture, click on . The page **Advanced settings** opens. For more details regarding this page, refer to the last steps of the relevant smart architecture addition procedure in the section [Adding a DNS Smart Architecture](#).
8. Click on  to complete the operation. The report opens and closes. If your configuration is still managing DNS servers, you can display them in the list **All servers** using the  button in the upper right corner.

To delete Amazon Route 53 NS resource records

1. In the sidebar, go to  **DNS** > **RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the column **Type**, filter the list to only display *NS* records.
4. In the field **Value**, filter the list to display *awsdns* records.
5. Tick the record(s) you want to delete.
6. In the menu, click on  **Delete**. The wizard **Delete** opens.
7. Click on  to complete the operation. The RR is marked  **Delayed delete** and is then no longer listed.

Managing Amazon Route 53 Zones

From the GUI you can manage the existing zones of an Amazon Route 53 server and add, edit or delete AWS zones as well.

However, there are some prerequisites, specificities and limitations to take into account.

Prerequisites

- The zone name must be unique: if several zones of the Amazon Route 53 server share the same name on your AWS account, you cannot manage any of them.

- The zone name must be FQDN and include a TLD.
- Make sure that SOLIDserver is using UTC system, otherwise any changes made from the GUI cannot be pushed to your AWS account server.

AWS Zones Specificities

- You can only add or import Master zones. For more details, refer to the sections [Adding a Master Zone](#) and [Importing Zones](#).
- Your zone contains by default an SOA specific to Amazon Route 53 and four *awsdns* NS records.
- On the zones properties page, the Name servers panel lists all the NS records of the zone.
- Adding, editing or deleting zones automatically refreshes the server and replicates your changes to the AWS account server.
- The report **Zones NS and IP addresses** allows you to retrieve the IP address of each of your NS records. For more details, refer to the section [DNS Zone Reports](#).

Limitations

- If your AWS zones contain records configured with the options *Health Check* and *Routing Policy*, they cannot be synchronized. These options are not supported by SOLIDserver. You must remove these options directly via your AWS account to then be able to synchronize and manage the zone via the GUI.

Managing Amazon Route 53 Records

From the GUI you can manage the existing records of your AWS zones, as well as add, edit or delete Amazon Route 53 resource records.

However, there are some prerequisites, specificities and limitations to take into account.

Prerequisites

- Make sure that SOLIDserver is using UTC system, otherwise, any change made from the GUI cannot be pushed to your AWS account server.

AWS Records Specificities

- Each zone contains by default four NS records with the value *awsdns*, specific to Amazon Route 53 servers.
- Each zone contains by default an SOA record named after one the NS records. For that reason, even if the server is managed by a smart architecture, the SOA name is not overwritten by the name of the smart architecture server.

Limitations

- Only a set of records is supported by Amazon Route 53 servers: NS, MX, A, AAAA, PTR, CNAME, TXT and SRV. For more details, refer to the section [Adding Resource Records](#).
- If your server is managed via a smart architecture, you can add other types of records, but their status is *N/A* and they are not taken into account, or replicated, by your Amazon Route 53 server.
- AWS records configured with the options *Health Check* and *Routing Policy* are not supported. If a zone contains them, it cannot be synchronized.

- The AWS zones resolution relies on the *awsdns* NS records, so you cannot delete them from a zone, an AWS server or a smart architecture managing an AWS server.

However, if you stop managing an Amazon Route 53 server via a smart architecture, you can remove the *awsdns* records from the smart architecture All zones list as they are no longer relevant.

Synchronizing DNS Servers

The synchronization of most servers is automatic but administrators can synchronize servers manually to integrate faster changes made to the zones or views databases.

Amazon Route 53 servers must be synchronized if they were edited directly from the AWS account.

Some data, like the panel *Sources* of physical servers, is only visible once the server has been successfully synchronized at least once.

To synchronize DNS servers

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Tick the zone(s) you want to synchronize.
3. In the menu, select **Edit > Synchronize**. The wizard **Synchronization** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The page reloads.

Editing DNS Servers

To edit any kind of DNS server configuration, you need to open its properties page and edit the relevant panel(s). Before proceeding, note that:

- The panels that do not contain the button *EDIT* cannot be edited.
- The basic server parameters are detailed in the addition procedures of each kind of physical server.
- For server managed via a smart architecture, some parameters can only be edited when you edit the smart. For instance, on EfficientIP servers managed via a smart, the boxes *Isolated* and *Use as DNSSEC resolver* and the advanced properties parameters are only available when you edit the smart.
- The SNMP protocol is no longer supported as managing protocol for a server, so **editing it automatically changes the management protocol to use SSL instead of SNMP**. This operation is **non-reversible**.
- EfficientIP DNS servers in version 7.x are managed through a dedicated service account using a randomly generated password. Note that:
 - A server can be managed by only one appliance. To switch the appliance managing the server, you need to edit your sever and type in your credentials again.
 - For servers added in previous versions, after an upgrade to version 7.0, to switch to this new management system, you need to edit the servers.
- EfficientIP DNS servers can only be edited if they are reachable, they must be up and running when you edit them.

For more details regarding all the server configuration possibilities (forwarding, recursion, transfer, blackhole, sortlist, etc.), refer to the chapter [Configuring DNS Servers](#).

You can edit EfficientIP, Nominum and Generic servers to secure their data exchanges with SOLIDserver. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).

To edit a DNS server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. At the end of the line of the server of your choice, click on **ⓘ**. The properties page opens.
4. Open all the panels using **☰**.
5. In the panel of your choice, click on **EDIT**. The corresponding wizard opens.
6. Make all the changes you need. For an EfficientIP server, from the panel **Main Properties** you can:
 - a. Tick the box **Use DNS as DNSSEC resolver**, to enable DNSSEC resolution on the server. For more details, refer to the chapter [Managing DNSSEC on Recursive Servers](#).
 - b. Tick the box **Configure enrollment parameters**, to manage a server that is already managed by another appliance or to switch to the new management system.

In the field "**Admin**" **account password**, enter your SSH password.
 - c. Tick the box **Enable HSM**, to start sign Master zones using HSM. For more details, refer to the chapter [HSM](#).
7. Click on **NEXT**, if need be, until you get to the last page of the wizard.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and refreshes.

Securing the Management of DNS Servers Using a TSIG Key

You can use TSIG keys to secure all the data exchanges between EfficientIP, Nominum and Generic servers and SOLIDserver. They are not supported by Microsoft servers.

By default, EfficientIP servers provide TSIG keys on the properties page of each DNS physical server. You must edit the server main properties to specify the existing TSIG key you want to use and secure data exchanges.

If you want to add the TSIG key, refer to the section [Configuring DNS Keys](#).

Note that once added on the properties page of a server:

- TSIG keys can be used in any of the server statements or in the statements of its views and zones.

At zone level you can set up dynamic update if you use the TSIG key specified on the server in the statement *allow-update*, for more details refer to the section [Configuring DNS Update Authorizations on a Zone](#).

- TSIG keys can be added to the ACL *admin* of the server. This allows to automatically secure the statement *allow-transfer* that, by default, grants access to the ACL *admin*. For more details, refer to the section [Configuring Access Control Lists For a Server](#).

If you manage your physical servers from a smart architecture, the TSIG keys of the smart architecture are pushed to the properties of each of the physical servers it manages. So keep in mind that **a TSIG key must be unique to each server**, you cannot use the same for several servers.

To select a TSIG key for EfficientIP and Nominum servers

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Right-click over the name of the EfficientIP DNS, EfficientIP DNS Package and Nominum server of your choice, click on . The wizard **Edit a DNS server** opens.
3. Click on **NEXT** until you get to the last page of the wizard.
4. Tick the box **Configure TSIG parameters** if it is not already ticked.
5. In the drop-down list **TSIG key name**, select the key of your choice.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

To configure a TSIG key for Generic servers

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Generic server of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
4. Click on **NEXT** until you get to the last page of the wizard.
5. Tick the box **Configure TSIG parameters** if it is not already ticked.
6. In the field **TSIG key name**, specify the name of the key.
7. In the drop-down list **TSIG key method**, select the method that suits your needs⁸. If you are not using an access key for this server, select *None*.
8. In the field **TSIG key value**, specify your key value.
9. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Deleting DNS Servers

At any time you can delete a server from the page *All servers* and stop managing it from SOLIDserver.

Keep in mind that **you cannot delete a physical server if:**

- It is managed by a smart architecture. For more details, refer to chapter [Managing DNS Smart Architectures](#).
- It is associated with an application. For more details, refer to chapter [Managing Applications](#).

⁸The standardized protocol for key codes is HMAC-MD5.

To delete a DNS server

1. In the sidebar, go to  **DNS** > **Servers**. The page **All servers** opens.
2. Filter the list if need be.
3. Tick the server(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The server might be marked  *Delayed delete* until it is no longer listed.

Defining a DNS Server as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a server as one of the resources of a specific group allows the users of that group to manage the server in question as long as they have the corresponding rights granted.

Granting access to a server as a resource also makes every item it contains available. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 36. Configuring DNS Servers

This chapter details the possible server configurations on its properties page, whether smart or physical. Any configuration set at server level is inherited by all the views and zones of the server.

Most options provide ACL configuration fields, keep in mind that in these fields the order of the elements listed is important: each restriction or permission is reviewed following the order set in the list.

If you configure any of these options at view or zone level, the value set at server level is overridden.

Configuring DNS Forwarding at Server Level

A forwarder is a DNS server that is designated to facilitate forwarding of queries for other DNS servers. It can manage name resolution for names outside of your network, such as names on the Internet, and improve the efficiency of name resolution for the computers in your network. Forwarding is used only for queries for which the server is not authoritative and does not have the answer in its cache.

Setting a DNS server as a forwarder allows to prevent leaving DNS information exposed outside of a network as your DNS servers do not need to send queries outside to their root hints. In addition, it allows to minimize the volume of external traffic which can be costly and inefficient for a network with a slow Internet connection or a company with high Internet service costs.

If you specify a list of forwarders on a smart server, you can set the forwarding to:

- **first:** the DNS server queries the forwarders first, if none of the forwarders in the list are responsive, the server looks for the answer itself.
- **only:** the DNS only forward queries to the forwarders in order to avoid an answer seeking.
- **none:** the forwarding is disabled. This is the default value.

SOLIDserver always sends the query to the forwarder with the lowest *round trip time* (RTT) in the list of forwarders configured. The RTT measures how long a remote name server takes to respond to queries. Each time an EfficientIP DNS server sends a query to a forwarder it starts an internal clock, and stops it when it receives a response. The RTT is stored to ensure that queries are sent to the proper forwarder.

Keep in mind that:

- The forwarding configuration set on a smart architecture is automatically inherited by the servers it manages. You can override the configuration directly on a physical server.
- **Any configuration set at view or zone level overrides the server level configuration.** For more details, refer to the sections [Configuring DNS Forwarding at View Level](#) and [Configuring DNS Forwarding at Zone Level](#).

To configure forwarding on a server or smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.

2. At the end of the line of the server or smart architecture of your choice, click on . The properties page opens.
3. Open the panel **Forwarding** using .
4. Click on **EDIT**. The wizard **Forwarding configuration** opens.
5. In the field **Forward mode**, select the mode of your choice according to the table below.

Table 36.1. Forward mode options at smart architecture or server level

Option	Description
None	The server does not send the queries to a forwarder and looks for the answer itself. This option is set by default. Selecting this option clears the list <i>Forwarders</i> . You cannot set <i>None</i> on a server managed by a smart architecture.
First	The server sends the queries to the forwarder(s) and, if it does not receive any answer, attempts to find an answer on its own.
Only	The server only forwards queries to the forwarder(s) listed.

6. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.
7. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displayed the new settings. In the panel **Forwarding**, your configuration is displayed.

You can set a specific forwarding configuration for physical servers managed via a smart architecture already configured with forward options. This new configuration is inherited by the views, zones and records of the physical server. Keep in mind that:

- When a forward mode is set on a smart architecture, you cannot set the forward mode to *None* on any physical server it manages. You can only set a different forward mode.
- Any configuration set at view or zone level overrides the server level configuration.

To configure a specific forward mode on a server managed via a smart architecture

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. Make sure the servers managed by your smart architectures are displayed. If not, on the right-end side of the menu, click on .
3. At the end of the line of the server or smart architecture of your choice, click on . The properties page opens.
4. Open the panel **Forwarding** using .
5. Click on **EDIT**. The wizard **Forwarding configuration** opens.
6. Tick the box **Overwrite the smart settings**. The page refreshes and displays additional fields.

7. In the field **Forward mode**, select *First* or *Only*. For more details, refer to the table [Forward mode options](#).

You cannot set the forward mode of a physical server managed via a smart architecture to *None*.

8. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

9. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displays the new settings. In the panel **Forwarding**, the *Forward* value is preceded by the message *Smart configuration is overwritten*.

To revert the specific configuration and inherit it again, edit the *Forwarding* to untick the box *Overwrite the smart settings*.

Configuring DNS Recursion at Server Level

In principle, authoritative name servers are sufficient for the operation of the Internet. However, with only authoritative name servers operating, every DNS query must start with recursive queries at the root zone of the DNS and each user system must implement resolver software capable of recursive operations. To improve performance, recursive servers cache the results of the lookups they perform. The processes of recursion and caching are intimately connected, then the terms recursive server and caching server are often used synonymously. The length of time for which a record may be retained in the cache of a caching name server is controlled by the field *Time To Live (TTL)* of each resource record. Typically, such caching servers, also called DNS caches, also implement the recursive algorithm necessary to resolve a given name starting with the DNS root through to the authoritative name servers of the queried domain. By default the DNS recursion function is enabled in SOLIDserver DNS.

A recursive query requires the DNS server to return requested DNS data, or locate the data through queries to remote DNS servers. When a DNS server receives a query for DNS data it does not have, it first sends a query to any specified forwarders. If a forwarder does not respond with any return, it resends the same query to the next configured forwarder until it receives an answer. If it receives no answer or a negative answer, then it sends a non-recursive query to specified internal root servers. If no internal root servers are configured, the DNS server sends a non-recursive query to the Internet root servers.

Enabling and Disabling the Recursion

If the recursion is enabled, the server always provides recursive query behavior if requested by the client. If it is disabled, the server only provides iterative query behavior - normally resulting in a referral. If the answer to the query already exists in the cache, it is returned irrespective of the value of this statement. This statement essentially controls caching behavior in the server.

By default, the DNS recursion is enabled. The DNS properties page displays the panel *Recursion* that allows you can set different DNS recursion configurations.

Keep in mind that **any configuration set at view or zone level overrides the server level configuration.**

To enable the DNS recursion

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
3. Open the panel **Recursion** using **☑**. If the **Recursion** is set to *No*, click on **EDIT**. The wizard **Recursion configuration** opens.
4. In the drop-down list **Recursion**, select *Yes*.
5. In the field **Recursive-clients**, type in the number of clients that you want to serve recursively. If you leave the field empty, the default value, *1000*, is applied. For more details, refer to the section [Limiting the Number of Clients Served Recursively](#).
6. Click on **NEXT**. The page **Allow recursion** opens. For more details, refer to the section [Limiting the Recursion at Server Level](#).
7. Click on **OK** to complete the operation.

To disable the DNS recursion

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
3. Open the panel **Recursion** using **☑**. If the **Recursion** is set to *Yes*, click on **EDIT**. The wizard **Recursion configuration** opens.
4. In the drop-down list, select *No*.
5. Click on **NEXT**. The page **Allow recursion** opens.
6. Click on **OK** to terminate the recursion disabling.

Limiting the Recursion at Server Level

By default, the EfficientIP DNS is allowed to serve all clients that send recursive queries. You can restrict it by defining a match list defining IP address(es) which are allowed to issue recursive queries to the server. Limiting the recursion allows to specify which hosts are allowed to make recursive queries through the DNS server. If the restriction of the recursion (*allow-recursion*) is not set then the restriction of caching (*allow-query-cache*) is applied if set, otherwise the restriction of queries (*allow-query*) is used if set, or the default (localnets; localhost;) is used. If the answer to the query already exists in the cache, it is returned irrespective of this statement.

Keep in mind that **any configuration set at view or zone level overrides the server level configuration.**

To set an allow-recursion match list at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⚙**. The properties page opens.

3. Open the panel **Recursion** using and click on **[EDIT]**. The wizard **Recursion configuration** opens.
4. Click on **[NEXT]**. The page **Allow recursion** opens.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.2. Restriction and permission parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **[ADD]**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons and .

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **[UPDATE]** or click on **[DELETE]** to remove an entry from the list. If you made changes that you do not want to save, click on **[CANCEL]**.

5. Click on **[OK]** to complete the operation. The report opens and closes. The properties page is visible again.

Limiting the Number of Clients Served Recursively

The statement *recursive-clients* allows to define the number of simultaneous recursive lookups the server performs on behalf of its clients. In other words, it allows you to set or limit the number of clients that your BIND server serves at the same time.

Once you set the recursion to *yes*, the *recursive-client* statement is enabled. To disable the *recursive-clients* statement, you must disable the recursion.

The statement default value is 1000, meaning that 1000 simultaneous lookup requests can be answered by the server. The minimum value is 1 and the maximum value is 4294967295.

You can set this statement on a smart architecture or on a physical server.

To limit the number of clients served recursively

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on . The properties page opens.

3. Open the panel **Recursion** using . If the **Recursion** is set to *No*, click on . The wizard **Recursion configuration** opens.
4. In the drop-down list, select *Yes*.
5. In the field **Recursive-clients**, the default value *1000* is displayed. Edit the value if need be, the accepted values are between 1 and 4294967295.
6. Click on . The page **Allow recursion** opens. For more details regarding the recursion configuration, refer to the section [Limiting the Recursion at Server Level](#) above.
7. Click on to complete the operation.

Configuring DNS Notify Messages at Server Level

The DNS notification promotes consistency between primary and secondary servers as it allows to notify slave zones of changes performed on the master zone. Configuring the Notify at server level allows to set the changes notification once, for all the master zones managed by the primary server. It obviously implies that this *primary* server contains master zones already configured with corresponding slave zones on the *secondary* server¹. Once the notification is sent to slave zones, the administrator decides if a zone transfer is relevant. For more details, refer to the section [Limiting Zone Transfers at Server Level](#).

Within SOLIDserver, the notification configuration is done from the panel *Notify* of the properties page. This panel displays:

- The type of notification configured for the server. You can set the **Notify** to *Yes* or *Explicit* or *No*.
- The **Also notify** statement IP address and port. This statement allows to notify the managing smart server of any slave zones updates.

Keep in mind that by default the also notify statement is unset. This implies that the smart server is informed of the changes performed on slave zones when it refreshes, every hour, and not instantly.

- The **Allow notify** directive of the server slave zones. For instance, you can allow all the servers of a network to notify the slave zones of your server or only a few.

Note that there is an implicit allow-notify directive set when you add a slave zone: when you set the Master IP address of the slave zone you are allowing the master zones of this server to send notify messages to your slave zone.

Keep in mind that **any configuration set at view or zone level overrides the server level configuration**.

To configure notify messages at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on . The properties page opens.
3. Open the panel **Notify** using and click on . The wizard **Notifying configuration** opens.

¹In this paragraph, to simplify the explanation, we work on the assumption that one server, the master server, contains only master zones and another, the secondary one, contains only slave zones. It is evidently not accurate: usually a server would manage both master and slave zones. However, it is customary to configure corresponding slave and master zones that are managed by a different server.

4. In the drop-down list **Notify**, configure the server notification behavior following the table below.

Table 36.3. DNS server notify types

Field	Description
No	No notify message is sent when changes are performed in the master zones.
Yes	The notify messages are sent to the target of the NS records of the master zone. They are also sent to the IP address(es) specified in the field <i>IP address</i> below.
Explicit	The notify messages are only sent to the IP address(es) specified in the field <i>IP address</i> below.

5. If you selected *Yes* or *Explicit*, you can set the IP address and port of the server(s) which slave zones should receive the messages:
 - a. In the field **IP address**, type in the IP address of another server. The notify messages are sent if you chose the notify type *Yes* or *Explicit*.
 - b. In the field **Port**, you can specify the port number that should receive the notify messages on the server you specified in the previous field.
 - c. Click on **ADD**. The IP address and port number are displayed in the list **Also notify** as follows: *<ip-address> port: <port-number>*. Repeat these actions for as many servers as needed.

You can edit the content of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

6. Click on **NEXT**. The page **Allow notify** opens. It allows to specify if the server slave zones can receive master zones notification messages.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.4. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons  and .

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again. Your configurations are displayed in the panel **Notify**.

Restricting DNS Queries at Server Level

The DNS queries can be restricted through the options `allow-query` and `allow-query-cache`. They both set an ACL list for IP addresses and/or network addresses, so keep in mind that **the order of the elements listed in the field ACL values is important** as each restriction or permission is reviewed following the order you set in the list.

Allow query

You can specify which hosts are allowed to issue DNS queries. The allow query properties can be configured for an entire server including all the zones it contains. By default, queries are allowed from the local host (`localhost`) and the local networks (`localnets`).

Keep in mind that **any configuration set at view or zone level overrides the server level configuration**.

To set an allow-query match list at server level

You can apply the procedure below, at zone level as well.

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⊞**. The properties page opens.
3. Open the panel **Access control** using **⊞**. This panel displays different options: **Allow-query**, **Allow-query-cache**, **Allow-transfer** and **Blackhole**.
4. Click on **EDIT** to change the configuration. The wizard opens.
5. On the page **Allow query**, set up the allow query match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.5. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <code><ip-address>/<prefix></code> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **[ADD]**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **[↕]** and **[↔]**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **[UPDATE]** or click on **[DELETE]** to remove an entry from the list. If you made changes that you do not want to save, click on **[CANCEL]**.

6. Click on **[NEXT]** twice to skip the pages **Allow-query-cache** and **Allow transfer**.
7. On the page **Blackhole**, click on **[OK]** to complete the operation. The report opens and closes. The properties page is visible again.

Allow query cache

You can set a list of the IP addresses that are allowed to issue queries on the local cache. The allow-query-cache properties are configured at server level and apply to the zones managed through the server.

Allow-query-cache statement particularities

The allow-query-cache is independent from the allow-query statement but closely linked to the allow-recursion statement.

If the recursion is set to no, the cache cannot be queried, so it is useless to set an allow-query-cache match list.

If the recursion is set to yes and the allow-recursion statement is not defined, by default the localhost and localnets are permitted to query the server cache.

If the recursion is set to yes and the allow-recursion statement is defined with a specific match list, the local cache access is granted to all the entries of the allow-recursion match list.

The match list defined controls recursive behavior as recursive queries would be useless without access to the local cache. Typically, if a host is in the allow-recursion match list, it could access the server the first time and get query result. However, if it is not part of the allow-query-cache match list then it would not be able to make the same query a second time as it would be saved on the cache to which it does not have access. On the contrary, if a host is in the allow-query-cache match list but not in the allow-recursion match list, it would only get results for queries already sent by another host with the proper access rights. Hence the need to configure carefully both these statements to avoid conflicts and absurd access configurations.

Keep in mind that **any configuration set at view level overrides the server level configuration**.

To set an allow-query-cache match list at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **[ⓘ]**. The properties page opens.
3. Open the panel **Access control** using **[☰]**. This panel displays different options: **Allow-query**, **Allow-query-cache**, **Allow-transfer** and **Blackhole**.
4. Click on **[EDIT]** to change the configuration. The wizard opens.
5. Click on **[NEXT]** to skip the page **Allow-query**.
6. On the page **Allow-query-cache**, set up the allow query cache match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.6. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **NEXT** twice to skip the page **Allow-transfer** and open the page **Blackhole**.
- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Limiting Zone Transfers at Server Level

DNS zone transfer is a type of DNS transaction employed to replicate and synchronize all copies of the zone used at each server configured to host the zone. SOLIDserver denies zone transfers by default to all DNS server. SOLIDserver supports the allow-transfer server option that allows to specify which hosts, networks, or TSIG keys are granted or denied the right to do transfers for all the zones it maintains.

Keep in mind that **any configuration set at view or zone level overrides the server level configuration**.

To set an allow-transfer match list at server level

You can also apply the procedure below at zone level.

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the server of your choice, click on **ⓘ**. The properties page opens.
- Open the panel **Access control** using **⊞**. This panel displays different options: **Allow-query**, **Allow-query-cache**, **Allow-transfer** and **Blackhole**.
- Click on **EDIT** to change the configuration. The wizard opens, each page corresponds to an option.

5. Click on **NEXT** twice to skip the pages **Allow-query** and **Allow query cache**.
6. On the page **Allow-transfer**, setup the transfer match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.7. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

7. Click on **NEXT**. The page **Blackhole** opens.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Configuring a Blackhole

You can set a list of the IP addresses and network addresses you consider as spam. The blackhole properties can be configured for an entire server including all the zones it contains. By default, queries are allowed from the local host and the local networks: all the addresses listed in the list cannot receive any response from the server or zones. The queries remain unanswered, in other words, ignored.

To set a blackhole match list at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⊞**. The properties page opens.
3. Open the panel **Access control** using **☐**. This panel displays different options: **Allow-query**, **Allow-query-cache**, **Allow-transfer** and **Blackhole**.
4. Click on **EDIT** to change the configuration. The wizard opens.
5. Click on **NEXT** to skip the pages **Allow-query**, **Allow query cache** and **Allow-transfer**.

- On the page **Blackhole**, set up the restrictions. You can deny access to networks and IP addresses, they can all be listed in the list **ACL values**. The table below details the available options of the field **Type**:

Table 36.8. Blackhole parameters

Type	Description
Network address	Deny query responses to an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Deny query responses to an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in an IPv4 or IPv6 address.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Configuring Client Resolver Cache Options at Server Level

From the properties page of an EfficientIP DNS server using the SSL protocol, you can edit the two options dedicated to client resolver cache memory via the panel *Options*:

lame-ttl

This option defines the amount of time a client should keep in its cache the information sent by a lame server that has been queried directly. It allows to limit the time the information is kept as, coming from a lame server, it might not be up-to-date and therefore potentially erroneous.

max-cache-size

This option limits the size of the cache memory of a server or view. When the cache memory size reaches this threshold, the server causes records to expire prematurely. The value 0 can be set to purge the cache only when the records TTL expires.

These options can be set at server or view level. For more details regarding the configuration on views, refer to the section [Configuring Client Resolver Cache Options at View Level](#). Keep in mind that **any configuration set at view level overrides the server level configuration**.

To set the option lame-ttl at server level

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
- Open the panel **Options** using **☐** and click on **EDIT**. The wizard **Options configuration** opens.
- In the field **Lame-ttl**, type in the value of your choice. This value is in seconds can be set between 30 and 1800. The default value is 600, the maximum value is 1800 seconds.
- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

To set the option max-cache-size at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Max-cache-size**, type in the value of your choice to set the cache memory size. This value is in bytes. The default value is *100m*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

Configuring EDNS Options at Server Level

The Extension mechanisms for DNS allow to add information to DNS messages and therefore send out larger packages or packages containing more parameters. The EDNS, also known as *EDNS (0)*, has been defined in RFC 6891.

If you configured DNSSEC on your server or are managing records that relay IPv6 information, we strongly recommend configuring EDNS: in both cases, the messages sent out usually exceed 512 bytes.

Within SOLIDserver, two EDNS options can be configured at server and view level on EfficientIP DNS servers using the SSL protocol:

edns-udp-size

This option sets the EDNS UDP buffer size advertised by the server when querying a remote server. It is set in bytes and allows to **specify the size of the packets that you receive**. Typically, you would set this option to enable UDP answers to pass through broken firewalls that block fragmented packets and/or packets greater than 512 bytes. The value set for this option is a preference.

max-udp-size

This option sets the maximum EDNS UDP message size sent by the server. It is set in bytes and allows to **specify the maximum size of the packets that you send** to a remote server. Typically, this option would be set to enable UDP answers to pass through broken firewalls that block fragmented packets and/or packets greater than 512 bytes.

For more details regarding the configuration on views, refer to the section [Configuring EDNS Options at View Level](#). Keep in mind that **any configuration set at view level overrides the server level configuration**.

To set the option edns-udp-size

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Edns-udp-size**, type in the size of received packets of your choice. This value is in bytes, and must be set between 512 and 4096. The default value is *4096*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

To set the option `max-udp-size`

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Max-udp-size**, type in the maximum size of the packets you send. This value is in bytes and must be set between 512 and 4096. The default value is *4096*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

Improving the Server Performance

You can configure the statement *minimal-responses* at server level. This statement can be used to reduce the size of the outgoing data and improve performances.

Enabling *minimal-responses* allows the server to only add records to the authority and additional data sections of the response if they are specifically required by the protocol. For instance, these sections are included in the delegations and negative responses.

By default, the statement *minimal-responses* is set to *yes* on BIND servers.

This statement can be set on a physical server or on a smart architecture.

To set the `minimal-responses` statement at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the drop-down list **Minimal-responses**, select *yes* or *no* according to your needs. By default, the statement is set to *yes* on BIND servers.
5. Click on **NEXT**. The last page of the wizard opens.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays your **minimal-responses** configuration.

Configuring a Sortlist at Server Level

The option *sortlist* is actually a statement that allows to set a preferential response order for equal A resource records, forming an RRset. In other words, it modifies the response packet received by the client resolver. It allows to put an end to cyclic round-robin responses to queries for the IP networks of your choice. You can define as many *sortlist* statements as you want on EfficientIP DNS servers using the SSL protocol. For each network of client IP addresses, you can set the order of the records of an A RRset: this list can contain all of the A records of the RRset if you want. The server checks if the client resolver IP address matches the sortlist defined and modify its response accordingly.

Within the GUI, the statement configuration is closely linked to the statement syntax in the zone file. Here below is an **example of the sortlist statement syntax in the zone file**:

In a zone file, the statement would look as follows for the zone `many.example.com` :

```
// zone file example.com
$ORIGIN example.com.
many IN A 192.168.3.6
      IN A 192.168.4.5
      IN A 192.168.5.5
      IN A 10.2.4.5
      IN A 172.17.4.5
```

The client-side server has a sortlist statement, set as follows:

```
options {
    ....
    sortlist {
        { // 1st preference block start
            192.168.4/24; // 1st client IP selection matches any of these
            {10.2/16; // return any of these response IPs as 1st preference
              172.17.4/24; // 2nd preference
            };
        }; // end first block
        { // second preference block
            192.168.5/24; // 2nd client IP selection matches any of these
            {192.168.4/24; // return any of these response IPs as 1st preference
              172.18.4/24; // 2nd preference
              10.2/16; // 3rd preference
            };
        }; // end second block
    }; // end sortlist
};
```

As you can see after the client IP, the response preferences are defined one after the other and separated by a semi-colon.

Keep in mind that **any configuration set at view level overrides the server level configuration.**

To define a sortlist statement at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **ⓘ**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Client address**, type in the client IP address/subnet. It must be composed of an IPv4 address containing 1 to 4 bytes followed by the prefix: *<IP address>/<prefix>*.
5. In the field **Sort address**, type in a list of IP addresses or subnets followed by a semi-colon. These addresses correspond to the value of an A record of the RRset for which you create the sortlist. The statement respects the order in which you typed in the addresses. The value must respect the format *<IP address>/<prefix>;* even if you only type in one sort address.
6. Once both fields are filled, click on **ADD** to move the client and sort addresses to the field **Sortlist**. Both values are displayed as follows: *{<client-IP-address>/<prefix> {<sort-IP-address>/<prefix>;}*; . By default, this field is empty.
7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays your sortlist as follows: *{<client_address_field_value>; {<first_sort_address>;<second_sort_address>;<etc>;}*; . There is one sortlist per client address defined.

Limiting the Number of Responses of a Server

The management of the Response Rate Limiting (RRL) from the GUI helps you mitigate the DNS amplification attacks targeting a victim.

Amplification attacks are Distributed Denial of Service (DDoS) tactics in which an attacker uses the IP address of any computer to send high volumes of forged queries using an authoritative DNS server. The attacker queries are usually small-sized but designed to generate large responses, thus generating an amplified traffic toward the victim. Considering that the attacker mimics the query format of the server, the client only sees a large number of responses and cannot know if the response is real or malicious. This high volume of responses is likely to make the victim's computer overload and eventually collapse.

In light of the increasing number of DDoS attacks on the Internet, in 2012 by Paul Vixie and Vernon Schryver proposed RRL as a method for preventing a caching server from being used in a DNS Amplification attack. It allows to maintain the type of queries that have been made, that way the server keeps track of the queries made and can limit the number of responses returned without changing the nature of the DNS.

Prerequisites

To be able to use this option:

- You must have SOLIDserver in version 5.0.4 or greater.
- You must manage an EfficientIP DNS server in version 5.0.4 or greater.
- If necessary, install the EfficientIP DNS packages in version 5.0.4 or greater. For more details, refer to the section [Managing BIND DNS Servers](#).

Configuring the RRL Parameters

RRL has to be configured from a DNS server, or smart architecture, properties page.

If you set it on a smart architecture that manages different types of servers, it only applies to the relevant servers. The settings are ignored by all the servers that do not support it.

To limit the number of responses of a server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⋮**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **[EDIT]**. The wizard **Options configuration** opens.
4. Click on **[NEXT]**. The last page of the wizard opens.
5. In the field **Maximum number of responses per second**, type in the number of your choice. This field defines a threshold of responses which, once met, drops all the additional queries unless you set a *Slip* value.

The field default value is 0, meaning the option is disabled. You can set a value between 1 and 1000. **We recommend that you set the value at 10 or higher.**

6. In the field **Slip**, you can type in a number between 0 and 10. This option allows to add an extra bit to the server responses, the Truncated (TC) bit, and makes the requestor use TCP to resend the query once they receive the truncated response. The number specified in the field defines every how many queries the TC bit and TCP are forced.

If the field is empty or set to 0, all the similar queries are dropped. If set to 1, the TC is set in the response to every query. If set to 2, it is set in the response to every other query. If set to 3, it is set in the response to every third query, etc.

We recommend that you set the Slip value to 1 because it limits cache poisoning attacks - it is worthless for attackers performing volumetric or PPS DDoS attacks - and is less CPU consuming for flooded resolvers.

7. If you are configuring a BIND server, in the section **Log only** you can tick the box. It allows to prevent the rate limiting function from operating and only display in the logs what would have happened.

On NSD servers, ticking the box disables RRL altogether.

8. Click on to complete the operation. The report opens and closes. The properties page is visible again. The settings are visible in the panel **Options**.

Limitations

- RRL cannot be set at view or zone level, it can only be configured for a whole server.
- RRL can only be configured on EfficientIP DNS servers or smart architectures.
- If RRL settings are configured on a smart architecture managing servers that are not compatible with RRL, for these servers the RRL configuration is ignored and the Multi-Status column provides details regarding incompatibilities.
- RRL can be set on a BIND/NSD Hybrid server but the option *Log only* disables RRL on NSD servers.

Configuring DNS64

SOLIDserver provides, for EfficientIP and BIND servers, a panel on the properties page dedicated to configuring DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers, or DNS64.

This mechanism was described in RFC 6147 and allows to rely on a NAT64 server to synthesize IPv6 addresses and enable communication between IPv6-only clients and IPv4-only DNS servers and get a valid response to their queries. Which is why DNS64 mechanism is useless on its own.

SOLIDserver does not include a NAT64 server but provides a wizard to configure DNS64 with the settings that you configured on your own NAT64 server.

Once NAT64 is configured, you can configure DNS64 on the server of your choice. All the fields available in the GUI are optional, DNS64 can perform the synthesis using only the *address* and *prefix* of the NAT64. It relies on your NAT64 server IPv6 address to IPv4 address translation settings to synthesize AAAA records from A records. The synthesis is performed as follows:

1. DNS64 uses the NAT64 server address and prefix configuration as the start of the synthesized IPv6 address.
2. The target IPv4 address is appended to the address and prefix in hexadecimal form.
3. The suffix is appended at the end of the synthesized IP address to get a 128 bit address to send back to the client in the AAAA record.

When your DNS64 configuration is over, two reverse zones are automatically created. That way, even the reverse queries of IPv6 clients can be synthesized and answered.

The mechanism is completely transparent: the client does not know that the queried server only manages IPv4 related data as, after querying it, the client receives records that can be used in an IPv6 environment.

Prerequisites

- You must have SOLIDserver in version 5.0.0 or greater.
- You must manage:
 - An EfficientIP DNS server in version 5.0.0 or greater; or
 - A BIND server in version 9.8 or higher (EfficientIP BIND package).
- You must manage your server through SSL.
- You must have a NAT64 server configured on your network.
- You must have the NAT64 server address and prefix ready when you start configuring. Note that depending on your NAT64 server vendor, you might not be able to set the address and prefix of your choice. In this case you must use the ones provided by the vendor during the DNS64 configuration.

Limitations

- DNS64 can "break" DNSSEC. For more details, refer to the substatement [break-dnssec](#).
- DNS64 cannot be configured on Hybrid DNS engines.

DNS64 Supported Substatements

SOLIDserver supports the configuration of a set of DNS64 substatements listed below. By default all substatements are unset in the GUI as they are all optional, the fields are empty or set to *None*.

dns64-server

Allows you to specify the name of a server. The server you declare in that field should be the server for which you are configuring DNS64. Once the substatement is set, the server name is saved in the SOA of the reverse zones that DNS64 automatically creates, that way the server can be queried in reverse without further configuration.

Default value: N/A

dns64-contact

Allows you to specify the email address of the contact managing the server - declared in the substatement *dns64-server* - and the zones it contains. The contact is saved as well in the SOA of the reverse zones that DNS64 automatically creates. No need to specify a contact if you did not specify *dns64-server* substatement.

Default value: N/A

address

Allows you to specify the IPv6 address set on the NAT64 server following the syntax *<start-of-IPv6-address>::*. Depending on the NAT64 server of your network, you can either set the address of your choice or use the one already set on the NAT64 server. Without the *prefix*, filling this field is useless.

Default value: N/A

prefix

Allows you to specify the prefix set on the NAT64 server. It should be: *32 bits, 40 bits, 48 bits, 56 bits, 64 bits* or *96 bits*. Depending on the NAT64 server of your network, you can either set the prefix of your choice or use the one already set on the NAT64 server. No need to specify a prefix if you did not specify an *address*.

Default value: N/A

suffix

Allows you to specify a suffix following the syntax *::<end-of-IPv6-address>*. The suffix is appended to the synthesized IP address (composed as follows: *<address-field-value>:<queried-ipv4-address-in-its-hexadecimal-form>*).

Default value: *::*. No need to specify a suffix if you specified a *96 bits* prefix, BIND default value is automatically applied.

break-dnssec

Allows you to interact with DNSSEC. If you set it to *yes*, DNS64 synthesizes the responses authoritative and recursive queries even if the resulting records are considered invalid by the queried signed zones. Which is why it is considered to "break" DNSSEC. If set to *no*, the DNS64 synthesis is not performed if zones are signed with DNSSEC.

Default value: None.

recursive-only

Allows you to decide to use DNS64 synthesis in case of recursive queries. If set to *yes*, only the server responses to recursive queries are synthesized. If set to *no*, all the server responses are synthesized.

Default value: None.

client

Allows you to set or use existing ACLs to define which DNS clients get a synthesized response to their queries.

Default value: none. If you only specify an *address* and *prefix* this value is automatically applied.

mapped

Allows you to set or use existing ACLs to prevent DNS64 from synthesizing some IPv4 addresses of your network. Any IPv4 address listed in the ACLs is ignored by DNS64 and therefore never returned as an IPv6 address in the response.

Default value: any. If you only specify an *address* and *prefix* this value is automatically applied.

exclude

Allows you to set or use existing ACLs to exclude a list of IPv6 addresses. As DNS64 is dedicated to synthesizing IPv6 addresses from IPv4 addresses, if any zone contains AAAA records DNS64 ignores the zone. Therefore it ignores the A records it might contain. Setting up the *exclude* substatement allows to ignore the AAAA records of a zone, your *exclude* substatement ACL must contain all the IPv6 addresses declared in your AAAA records. That way you can make sure that DNS64 synthesizes all the A records of the zone. The IPv6 addresses you declare must be part of or match the range of addresses declared for the NAT64 server (*address* and *prefix*).

Default value: none. If you only specify an *address* and *prefix* this value is automatically applied.

Configuring DNS64 On a Server

You can configure DNS64 from the properties page of a physical server or a smart architecture.

To configure DNS64 on a server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
3. Open the panel **DNS64** using **⌵** and click on **[EDIT]**. The wizard **DNS64 configuration** opens.
4. Fill in the fields **DNS64 Server**, **DNS64 Contact** following the expected format detailed in the section [DNS64 Supported Substatements](#).
5. In the drop-down list **Prefix**, select *32 bits*, *40 bits*, *48 bits*, *56 bits*, *64 bits* or *96 bits*. For more details, refer to the section [DNS64 Supported Substatements](#).
6. Once you selected a Prefix, the field and drop-down lists **Address**, **Suffix**, **Break-dnssec** and **Recursive-only** appear. Fill them in according to your needs. For more details, refer to the section [DNS64 Supported Substatements](#).

Note that the field *Suffix* is not displayed if you select the Prefix *96 bits*.

You can choose to only configure the *Address* and *Prefix* for your DNS64 server and contact. In which case, SOLIDserver sets the substatement *client*, *mapped* and *exclude* to their [DNS64 Supported Substatements](#). To do so, click on **NEXT** until the last page of the wizard and click on **OK** to commit your configuration.

7. Click on **[NEXT]**. The page **DNS64 configuration: Client** opens.

You can define an ACL that lists the DNS clients that get a synthesized response to their queries. For more details regarding the substatement, refer to the section [DNS64 Supported Substatements](#).

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.9. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **[ADD]**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **⬅** and **➡**.

All the entries of the ACL values constitute your ACL.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

8. Click on **NEXT**. The page **DNS64 configuration: Mapped** opens.

You can define an ACL that lists the IPv4 addresses that are ignored by DNS64. The ACL configuration is detailed in the step 6 of this procedure. For more details regarding the sub-statement, refer to the section [DNS64 Supported Substatements](#).

9. Click on **NEXT**. The page **DNS64 configuration: Exclude** opens.

You can define an ACL that lists the IPv6 addresses that are ignored by DNS64. The ACL configuration is detailed in the step 6 of this procedure. For more details regarding the sub-statement, refer to the section [DNS64 Supported Substatements](#).

10. Click on **OK** to complete the operation. The report opens and closes. The properties page is available again, the panel **DNS64** displays your configuration.

To disable DNS64, you must open the wizard *DNS64 configuration*, empty all the fields and select *None* in the drop-down lists.

Configuring DNS Sources at Server Level

Configuring DNS source allows to set physical interfaces at server level that are systematically used for all notify operations and zone transfer. This information can only be set on EfficientIP DNS physical server using the SSL protocol and is inherited by the server views and zones and displayed accordingly on their properties page.

From the **Sources** and **Sources V6** panels, through their IP address, you can configure physical interfaces that should be used for the server transfer and notify options. These panels only appear after the first synchronization of the physical server. When editing these panels, you can define the following statements:

query-source

This statement allows to define the IPv4 address and/or port used as the source of the server or view outgoing queries. By default, BIND uses any server or view interface IP address and a random port for outgoing queries.

Using a fixed port number allows to control UDP operations but can be extremely dangerous: it can lead to cache poisoning if used with any caching DNS server definition as any attacker would need to guess the transaction ID to get both the specified interface IP address and port number. This statement is displayed on servers and views properties page.

query-source-v6

This statement allows to define the IPv6 address and/or port used as the source of the server or view outgoing queries. By default, BIND uses any server or view interface IP address and a random port for outgoing queries.

Using a fixed port number allows to control UDP operations but can be extremely dangerous: it can lead to cache poisoning if used with any caching DNS server definition as any attacker would need to guess the transaction ID to get both the specified interface IP address and port number. This statement is displayed on servers and views properties page.

transfer-source

This statement allows to determine the IPv4 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only

valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

transfer-source-v6

This statement allows to determine the IPv6 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

use-alt-transfer-source

This statement allows to set the use of an alternate interface IP address for the transfer if the *transfer-source* or the *transfer-source-v6* were to fail. This statement configuration is displayed on the physical server, view and slave zones properties page.

This statement definition is only configurable from the panel **Sources** but applies to interfaces whether they were identified through an IPv4 or an IPv6 address.

Its default value is *no* if the server contains views and *yes* if the server does not contain any view.

alt-transfer-source

This statement allows to determine the alternate IPv4 address of the interface used to execute the zones transfer on the server if the *transfer-source* fails and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

alt-transfer-source-v6

This statement allows to determine the alternate IPv6 address of the interface used to execute the zones transfer on the server if the *transfer-source-v6* failed and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

notify-source

This statement allows to define the IPv4 address of the physical interface used for all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

notify-source-v6

This statement allows to define the IPv6 address of the physical interface used all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

Keep in mind that **any configuration set at view or zone level overrides the server level configuration**.

To set IPv4 DNS sources at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on **ⓘ**. The properties page opens.
3. Open the panel **Sources** using **☰** and click on **[EDIT]**. The wizard **Configuration: Sources** opens.
4. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise, all the transfer operations would fail.

- a. In the field **Query-source address**, type in the IPv4 address of the interface used for outgoing queries.
- b. In the field **Query-source port**, you can type in the port number used for outgoing queries. Keep in mind that specifying a port number can lead to cache poisoning if DNS server definitions are not set properly.
- c. In the field **Transfer-source address**, type in the IPv4 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance.
- d. In the field **Transfer-source port**, you can specify which port on the interface should be used.
- e. In the drop-down list **Use-alt-transfer-source**, set the use of an alternate interface if need be.

Table 36.10. Use-alt-transfer-source parameters

Parameter	Description
none	This is the default value of the <i>use-alt-transfer-source</i> statement. If your server contains views, it corresponds to <i>no</i> . If your server does not contain any view, it corresponds to <i>yes</i> .
no	This value disables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails. Go to step 5 to set the notify-source statements related fields.
yes	This value enables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails. In this case, you need to set the alternate interface IP address (and port if you want) through the <i>alt-transfer-source</i> and <i>alt-transfer-source-v6</i> statements in the following steps.

The use-alt-transfer-source statement applies to the alternate interfaces declared through IPv4 address (Alt-transfer-source address) and IPv6 address (Alt-transfer-source address-v6).

- f. If you enabled the use of an alternate interface, in the field **Alt-transfer-source address**, type in the IPv4 address of the alternate interface. It must also be configured on the appliance.
 - g. If you enabled the use of an alternate interface, in the field **Alt-transfer-source port**, you can specify which port on the interface should be used.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise, all the notify operations would fail.
 - a. In the field **Notify-source address**, type in the IPv4 address to be used for outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source port**, you can specify which port on the interface should be used.
 6. Click on to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

To set IPv6 DNS sources at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on . The properties page opens.

3. Open the panel **Sources** using and click on . The wizard **Configuration: Sources** opens.
4. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the transfer operations would fail.
 - a. In the field **Query-source-v6 address**, type in the IPv6 address of the interface used for outgoing queries.
 - b. In the field **Query-source-v6 port**, you can type in the port number used for outgoing queries. Keep in mind that specifying a port number can lead to cache poisoning if DNS server definitions are not set properly.
 - c. In the field **Transfer-source-v6 address**, type in the IPv4 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance. If you defined the *use-alt-transfer-source* statement in the Sources panel, it applies to the alternate interfaces declared in IPv4 (*Alt-transfer-source address*) and IPv6 (*Alt-transfer-source address-v6*).
 - d. In the field **Transfer-source-v6 port**, you can specify which port on the interface should be used.
 - e. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 address**, type in the IPv6 address of the alternate interface. It must also be configured on the appliance.
 - f. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 port**, you can specify which port on the interface should be used.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise, all the notify operations would fail.
 - a. In the field **Notify-source-v6 address**, type in the IPv6 address to be used for the outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source-v6 port**, you can specify which port on the interface should be used.
6. Click on to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

Authenticating the Zones Dynamic Update from the Server

Within SOLIDserver, dynamic update relies on TSIG keys and can only be configured for the master zones if the server they belong to is properly configured and secured. Keep in mind that configuring dynamic update can be dangerous as any client could update the server data. You can configure dynamic update on all DNS servers except Amazon Route 53 DNS servers. For more details regarding dynamic update, refer to the chapter [Implementing Dynamic Update](#).

To set up dynamic update, you must:

1. Secure the server with a unique TSIG key. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
2. Configure the server for dynamic update:
 - Either edit the default ACL *admin* to include the same TSIG key to its permissions.

- Or create a new ACL that includes the TSIG key to its permissions. For more details, refer to the section [Configuring Access Control Lists For a Server](#).
3. Configure the statement *allow-update* of your master zones with the same TSIG key. For more details, refer to the section [Configuring DNS Update Authorizations on a Zone](#).

Note that if you edited the ACL *admin* of the server, the configuration is complete because, by default, the ACL *admin* of the physical server is specified in the statement *allow-update* of the master zones.

Note that instead of configuring the ACL *admin* to allow the securing TSIG key, you could allow the IP address of SOLIDserver and restrict the default permissions. However, **allowing updates based on the requestor IP address is insecure**, we strongly recommend using the TSIG key protocol filtering rather than an IP address based filtering.

To edit the ACL admin of a server to allow dynamic update

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on **⌵**. The properties page opens.
3. In the panel **ACL**, in the list **DNS ACLs** select *admin*.
4. Click on **EDIT**. The wizard **ACL configuration Edit a DNS server** opens.
5. Delete the ACL *any* to make sure that only SOLIDserver can edit the server and its content:
 - a. In the list **ACL values**, select *any*. The fields *Type*, *Restriction* and *ACL* are refreshed.
 - b. Click on **DELETE**. The ACL is no longer listed.
6. Specify your TSIG key:
 - a. In the drop-down list **Type**, select *TSIG key*.
 - b. In the drop-down list **Restriction**, select *Allow*.
 - c. In the field **Key**, select the same TSIG key than the one you used to secure the server. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
 - d. Click on **ADD**. The TSIG key is moved to the list **ACL values**.
7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.
8. Once the ACL *admin* is configured with the TSIG key, the master zones it contains are all configured for dynamic update as the statement *allow-update* grants by default access to this ACL.

To add an ACL that allows dynamic update

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
3. Open the panel **ACL** using **⌵**.
4. Click on **ADD**. The wizard **ACL configuration** opens.
5. In the field **ACL name**, name your ACL.

6. In the drop-down lists **Type**, select *TSIG key*.
7. In the drop-down lists **Restriction**, select *Allow*.
8. In the field **Key**, select the same TSIG key than the one you used to secure the server. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
9. Click on **ADD**. The TSIG key is moved to the list **ACL values**.
10. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.
11. Once the ACL is created, you must add it to the permissions of the statement *allow-update* of the master zone(s) of the server. For more details, refer to the section [Configuring DNS Update Authorizations on a Zone](#).

Configuring Access Control Lists For a Server

The Access Control List (ACL) is a match list that allows to grant or deny access to a network device, IP address, TSIG keys or even another ACL. For more details regarding TSIG keys, refer to the section [Configuring DNS Keys](#).

When set at server level, the ACLs can be used on the views and zones of the server. **Once created, you can use an ACL to configure the statements *allow-recursion*, *allow-notify*, *allow-query*, *allow-query-cache*, *allow-transfer*, *blackhole* at any of relevant level of the DNS hierarchy.** You could, for instance, create one ACL that specifies which part of the network is denied access or the IP address of the server that should always receive the notification messages, etc.

Note that by default EfficientIP servers provide the ACL *admin*. You can edit it according to your needs but you cannot delete it.

To create an ACL at server level

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **ⓘ**. The properties page opens.
3. Open the panel **ACL** using **⌵**.
4. Click on **ADD**. The wizard **ACL configuration** opens.
5. In the field **ACL name**, name your ACL.
6. Configure the ACL according to your needs.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 36.11. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at

Type	Description
	server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **↕** and **↕**.

All the entries of the ACL values constitute the content of your ACL.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again. Your ACL is listed in the panel **ACL**.

Once created, an ACL includes permissions and restrictions that you allow or deny access to depending on the configuration you set:

- If you allow access to the ACL, every permission it contains are granted access to, every restriction it contains are denied access to.**
- If you deny access to the ACL, the contrary is set: every permission it contains are denied access to, every restriction it contains are granted access to.**

You can add as many ACL as you want on a server.

To edit an ACL at server level

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
- Open the panel **ACL** using **⌵**.
- In the list **DNS ACLs**, select an ACL and click on **EDIT**. The wizard **ACL configuration** opens.
- To add permissions and restrictions to the ACL:
 - In the drop-down lists **Type**, select *Network address, IP address, ACL or TSIG key*.
 - In the drop-down lists **Restriction**, select *Allow or Deny*.
 - In the last drop-down list, specify or select the value of your choice. For more details, refer to the procedure [To create an ACL at server level](#).
 - Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!).
- To edit the permissions and restrictions of the ACL:

- a. In the list **ACL values**, select an entry. The fields above the list reload and display the current configuration.
 - b. Edit the fields according to your needs. If you made changes that you do not want to save, click on **CANCEL**.
 - c. When your changes are made, click on **UPDATE**. The fields reload, the entry is edited.
7. To change the order of the permissions and restrictions:
- a. In the list **ACL values**, select an entry.
 - b. Move it up or down the list using the buttons **▲** and **▼**.
8. To delete permissions and restrictions:
- a. In the list **ACL values**, select an entry. The fields above the list reload and display the current configuration.
 - b. Click on **DELETE**. The entry is no longer listed.
9. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

To delete an ACL at server level

1. Make sure the ACL is not used in any statement of the server, view or zone.
2. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
3. At the end of the line of the server of your choice, click on **⌵**. The properties page opens.
4. Open the panel **ACL** using **⌵**.
5. In the list **DNS ACLs**, select an ACL you want to delete.
6. Click on **DELETE**. The wizard **DELETE** opens.
7. Click on **OK** to complete the operation. The report opens and closes. The ACL is no longer listed.

Configuring DNS Keys

SOLIDserver supports the use of *Transaction Signatures* (TSIG) keys to encrypt and authenticate every DNS data exchange between SOLIDserver itself and your DNS servers or clients.

The information is encrypted via a technique called *HMAC* (Keyed-Hashing for Message Authentication, see RFC 2104) which employs a shared secret and a one-way cryptographic hash function to sign data. This shared secret is used as a password known only to the two parties involved in the exchange.

From the properties page of EfficientIP, Nominum and Generic servers as well as smart architectures you can add, edit and delete TSIG keys. Once a key is added, you can use it:

- To secure the server with a unique TSIG key. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
- In your ACLs at server, view and/or zone level. For more details, refer to the chapters [Configuring DNS Servers](#), [Configuring DNS Views](#) and [Configuring DNS Zones](#).

- When adding and editing slave zones, RPZ or not, and stub zones. For more details, refer to the chapters [Managing DNS Zones](#) and [DNS Firewall \(RPZ\)](#).
- To set up dynamic update for you master zones. For more details, refer to the sections [Configuring Access Control Lists For a Server](#) and [Authenticating the Zones Dynamic Update from the Server](#).

Note that TSIG keys are not supported by Microsoft servers. However, you can configure their zones for dynamic update via GSS-TSIG keys.

For more details, refer to the section [Implementing Dynamic Update](#).

To add a TSIG key

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⊞**. The properties page opens.
3. Open the panel **Keys** using **⊞** and click on **ADD**. The wizard **Add a DNS key** opens.
4. In the field **Key name**, name the key.
5. Click on **NEXT**. The wizard **TSIG Key configuration** opens.
6. A valid *HMAC-SHA512* key is automatically set in field **TSIG Key value**. If need be, edit the parameters described below:

Table 36.12. DNS key configuration parameters

Field	Description
Key name	The key name, is a string starting with a letter or underscore, followed by any number of letters, numbers, or underscores.
Key algorithm	The key algorithm can be either <i>hmac-sha512</i> , <i>hmac-sha384</i> , <i>hmac-sha256</i> , <i>hmac-sha224</i> , <i>hmac-sha1</i> or <i>hmac-md5 (obsolete)</i> .
TSIG Key value	The key value is the secret to be used by the algorithm, and is treated as a base-64 encoded string.

7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

To edit a TSIG key

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⊞**. The properties page opens.
3. Open the panel **Keys** using **⊞** and select the key you want to edit.
4. Click on **EDIT**. The wizard **TSIG Key configuration** opens.
5. In the field **TSIG Key value**, modify the data as needed.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

To delete a TSIG key

1. Make sure the key is no longer used in any statement of the server, view or zone.
2. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.

3. At the end of the line of the server of your choice, click on **■**. The properties page opens.
4. Open the panel **Keys** using **■**.
5. In the list **DNS keys**, select the key you want to delete in the list **DNS keys**.
6. Click on **DELETE**. The wizard **Delete** opens.
7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and the key is no longer listed.

Including Non-Supported DNS Settings

SOLIDserver provides the possibility to include parameters in BIND configuration files that are not supported by SOLIDserver.

Only administrators with a good understanding of BIND configuration files should perform these operations.

The prerequisites, limitations and procedures to follow are all described in the appendix [Configuring Non-Supported BIND Options](#).

Once the non-supported options are included to the configuration file, they are processed like any other option unless their syntax is incorrect or the option itself was already set. In which case, the included options are ignored until the proper changes have been made.

Configuring Anycast DNS

DNS anycast is a methodology useful if your deployment includes multiple geographically distributed sites. It improves the service high availability and reliability by improving the redundancy of the DNS service. Your DNS clients always query the same IP address(es) but their packets are systematically routed to the nearest server in the topology. The term "nearest" does not apply to the servers geographical repartition: if the closest server is down, the clients are redirected to the nearest running server in the topology. This avoids using remote servers based on the IP address alone and ensures that DNS clients are querying their local servers first. It can be implemented on recursive and authoritative DNS servers. In addition to sharing the workload, this configuration helps mitigating a DDoS attack by diluting its effects

You can implement anycast via the routing protocols OSPF, BGP and IS-IS and relies on a Quagga package, a host-based routing software, already stored on the appliance.

No matter the protocol chosen keep in mind that:

- Once anycast is implemented, the routers are able to redirect clients to the nearest server if need be.
- The Quagga configuration is automatically saved in the appliance backup file.

Implementing Anycast Using OSPF

OSPF is an intra network protocol. Successfully implementing anycast using OSPF on your network requires:

1. Meeting the [Prerequisites](#) to implement anycast.
2. [Configuring the Appliance for OSPF Anycast](#) to make sure it uses the Quagga package that enables anycast.

3. [Configuring the Quagga Package for OSPF Routing](#) to set the configuration that suits your needs.
4. [Making Sure DNS Anycast Was Properly Configured for OSPF](#) once the configuration is complete.

Prerequisites

- Several servers in a pool must share 1 or anycast IPs.
- The servers must advertise their IP(s) to their neighboring routers.
- The routers exchange the routes information. That way if one server fails, the routers automatically recompile the routing tables to redirect the DNS clients.
- The 3 steps anycast configuration must be completed on all the appliances that manage a DNS server that you intend to include to the anycast routing scheme. This applies whether the servers are managed via a smart architecture or not.

With this type of topology, the anycast IP address is advertised from multiple locations and the router ends up choosing the best path to that IP address, according to the metric in use by the routing protocol. Once you finished the configuration detailed in the sections below, the DNS servers managed via SOLIDserver use anycast.

Configuring the Appliance for OSPF Anycast

SOLIDserver contains a Quagga package that must be taken into account in the system configuration file to be used.

To successfully configure the package you must:

1. Meet the [Prerequisites](#).
2. Edit the `rc.conf` file to make sure it takes into account the package.
3. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*`.
4. Reboot the appliance. After the reboot, these files `running_conf.cf*` and `previous_conf.cf*` are created again and take into account the changes.
5. Configure the Quagga package. For more details, refer to the section [Configuring the Quagga Package for OSPF Routing](#).

To configure the appliance for anycast DNS

1. Edit the system configuration file.
 - a. Open a shell session and connect to your appliance with `root` credentials.
 - b. Open the file `/etc/rc.conf` to edit it.
 - c. Enable Quagga and make sure the file is configured as follows:

```
quagga_daemons="zebra ospfd"2
quagga_enable="YES"
```

- d. Add the following line to the file to specify the anycast dedicated IP address:

```
# Expected syntax: ifconfig_<interface-name>_alias0="inet <IP-address> netmask <netmask>"
# Example:
```

²To configure the appliance for both OSPF and BGP, you can specify the following: `quagga_daemons="zebra bgpd ospfd"` .

```
ifconfig_<interface-name>_alias0="inet <IP-address> netmask <netmask>"
```

e. Save your changes.

2. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*` using the command:

```
rm /tmp/running_conf.cf* /tmp/previous_conf.cf*
```

3. Reboot the appliance using the command:

```
reboot
```

Now you need to configure the package following the section below.

Configuring the Quagga Package for OSPF Routing

Once you configured the appliance to take into account the Quagga package, as detailed in the section [Configuring the Appliance for OSPF Anycast](#), you can configure the package and OSPF routing.

The package configuration implies:

1. Making sure that the firewall rule 36 using the OSPF protocol is enabled. Basically, this ensures that anycast management traffic and inbound messages are allowed.
2. Creating the Quagga and OSPF dedicated configuration files.
3. Restarting Quagga.
4. Checking the logs.

To make sure the anycast dedicated firewall rule is enabled

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Firewall rules**. The page **Firewall rules** opens.
3. In the column **Protocol**, type in `ospf` and hit Enter. The rule is the only one listed.
4. In the column **Action**, make sure it is marked `allow`.

To create the quagga dedicated configuration files

1. Open a shell session and connect to your appliance with `root` credentials.
2. Go to the directory `/data1/etc/quagga`.
3. In this directory, create the zebra configuration file using the following commands:

```
# emacs zebra.conf
```

It should contain the appliance hostname, administrator passwords, anycast IP address, anycast VIP(s) address and log file location like in the example below.

```
# more /data1/etc/quagga/zebra.conf | grep -v \!  
hostname dns-anycast-1  
password mypassword  
enable password mypassword  
  
# Specify the name of the interfaces used by your clients. "bgel" and "lo0" are examples.  
interface bgel  
ip address 192.168.53.2/24
```

```
interface lo0
ip address 192.168.55.2/32

log syslog debugging
log facility syslog
```

4. In this directory, create the OSPF configuration file using the following commands:

```
# emacs ospfd.conf
```

It should contain the appliance hostname, authentication details, response time, interfaces dedicated to OSPF, access list and log file location like in the example below.

```
# more /data1/etc/quagga/ospfd.conf | grep -v \!
hostname dns-anycast-1

# Specify the proper interface. "bgel" was specified in zebra.conf in our example.
interface bgel
ip ospf authentication message-digest
ip ospf message-digest-key 1 md5 mypassword
ip ospf priority 0
ip ospf hello-interval 1
ip ospf dead-interval 5

router ospf
log-adjacency-changes
ospf router-id 192.168.53.2
area 20 authentication message-digest
area 20 nssa
network 192.168.53.0/24 area 20
redistribute connected metric-type 1
distribute-list ANYCAST out connected
!
access-list ANYCAST permit 192.168.55.2/32

log syslog debugging
log facility syslog
```

To restart quagga

1. Open a shell session and connect to your appliance with *root* credentials.
2. Check the Quagga status using the following command:

```
/usr/local/etc/rc.d/quagga status
```

3. Restart Quagga using the following command:

```
/usr/local/etc/rc.d/quagga restart
```

To check quagga log file

1. Open a shell session and connect to your appliance with *root* credentials.
2. In the file `/var/log/zebra.log` you can check the Quagga dedicated logs. If everything went well, you should have three lines similar to the ones below:

```
// example of a successful configuration

Feb 25 09:46:02 dns1-anycast ospfd[18600]: Packet[DD]: Neighbor 192.168.53.1 Negotiation done
(Master).
Feb 25 09:46:02 dns1-anycast ospfd[18600]: AdjChg: Nbr 192.168.53.1 on bgel:192.168.53.2: Loading
-> Full (LoadingDone)
Feb 25 09:46:02 dns1-anycast ospfd[18600]: nsm_change_state(192.168.53.1, Loading -> Full):
scheduling new router-LSA origination
```

Making Sure DNS Anycast Was Properly Configured for OSPF

You can make sure that DNS anycast is successfully implemented, on the router itself and on SOLIDserver, via a set of commands that ensure that the IP addresses used during the configuration are part of the routes.

To display the OSPF routes on Cisco routers

1. Connect to SOLIDserver via a shell session.
2. Run the following command:

```
router>show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

    192.168.56.0/32 is subnetted 1 subnets
O N1 192.168.56.2 [110/21] via 192.168.53.2, 00:26:02, FastEthernet0/1
C    192.168.54.0/24 is directly connected, FastEthernet0/0
C    192.168.53.0/24 is directly connected, FastEthernet0/1
router>
```

To display the OSPF neighbors on Cisco routers

1. Connect to your Cisco router via a shell session.
2. Run the following command:

```
router>show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.53.2    1     FULL/BDR        00:00:38   192.168.53.2  FastEthernet0/0
```

To display the OSPF neighbors on SOLIDserver

1. Connect to SOLIDserver via a shell session.
2. Run the following command to connect to the zebra service:

```
# vtysh
Hello, this is Quagga (version 1.0.20160315).
Copyright 19962005 Kunihiro Ishiguro, et al.
dnslanycast#
```

3. Run the command

```
router>show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time   Address        Interface        RXmtL  RqstL  DBsmL
192.168.53.1    1     FULL/DR         36.591s    192.168.53.1  eth0:200.0.0.1  0      0      0
```

Implementing Anycast Using BGP

BGP is an inter network protocol that establishes communication between two AS (Autonomous System) numbers. Successfully using anycast for BGP routing protocol on your network requires:

1. Meeting the [Prerequisites](#) to implement anycast.

2. [Configuring the Appliance for BGP Anycast](#) to make sure it uses the Quagga package that enables anycast.
3. [Configuring the Quagga Package for BGP Routing](#) to set the configuration that suits your needs.
4. [Making Sure DNS Anycast Was Properly Configured for BGP](#) once the configuration is complete.

Prerequisites

Anycast implementation prerequisites:

- Several servers in a pool must share 1 or anycast IPs.
- The servers must advertise their IP(s) to their neighboring routers.
- The routers exchange the routes information. That way if one server fails, the routers automatically recompile the routing tables to redirect the DNS clients.
- The 3 steps anycast configuration must be completed on all the appliances that manage a DNS server that you intend to include to the anycast routing scheme. This applies whether the servers are managed via a smart architecture or not.

With this type of topology, the anycast IP address is advertised from multiple locations and the router ends up choosing the best path to that IP address, according to the metric in use by the routing protocol. Once you finished the configuration detailed in the sections below, the DNS servers managed via SOLIDserver use anycast.

BGP routing protocol prerequisites, you must have the following information ready:

- Make sure the network flows are properly configured, as detailed in the appendix [Matrices of Network Flows](#).
- Have the following information ready:
 - The local AS Number, EfficientIP AS Number.
 - The remote AS Number, the client AS Number.
 - The Anycast IP address.
 - The Neighbor IP address.

Configuring the Appliance for BGP Anycast

SOLIDserver contains a Quagga package that must be taken into account in the system configuration file to be used.

To successfully configure the package you must:

1. Meet the [Prerequisites](#).
2. Edit the rc.conf file to make sure it takes into account the package.
3. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*`.
4. Reboot the appliance. After the reboot, these files `running_conf.cf*` and `previous_conf.cf*` are created again and take into account the changes.
5. Configure the Quagga package. For more details, refer to the section [Configuring the Quagga Package for BGP Routing](#).

To configure the appliance for anycast DNS

1. Edit the system configuration file.
 - a. Open a shell session and connect to your appliance with *root* credentials.
 - b. Open the file `/etc/rc.conf` to edit it.
 - c. Enable Quagga and make sure the file is configured as follows:

```
quagga_daemons="zebra bgpd"3
quagga_enable="YES"
```

- d. Add the following line to the file to specify the anycast dedicated IP address:

```
# Expected syntax: ifconfig_<interface-name>_alias0="inet <IP-address> netmask <netmask>"
# Example:
ifconfig_bge1_alias0="inet 192.168.53.2 netmask 255.255.255.0"
```

Save your changes.

2. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*` using the command:

```
rm /tmp/running_conf.cf* /tmp/previous_conf.cf*
```

3. Reboot the appliance using the command:

```
reboot
```

Now you need to configure the package following the section below.

Configuring the Quagga Package for BGP Routing

Once you configured the appliance to take into account the Quagga package, as detailed in the section [Configuring the Appliance for BGP Anycast](#), you can configure the package for BGP routing.

The package configuration implies:

1. Creating the Quagga and BGP dedicated configuration files.
2. Restarting Quagga.

To create the quagga dedicated configuration files

1. Open a shell session and connect to your appliance with *root* credentials.
2. Go to the directory `/data1/etc/quagga`.
3. In this directory, create the zebra configuration file using the following commands:

```
# emacs zebra.conf
```

It should contain the appliance hostname, administrator passwords, anycast IP address, anycast VIP(s) address and log file location like in the example below.

```
# more /data1/etc/quagga/zebra.conf | grep -v \!
hostname dns-anycast-1
password mypassword
enable password mypassword
```

³To configure the appliance for both BGP and OSPF, you can specify the following: `quagga_daemons="zebra bgpd ospfd"`.

```
# Specify the name of the interfaces used by your clients. "bge1" and "lo0" are examples.
interface bge1
ip address 192.168.53.2/30
interface lo0
ip address 192.168.55.2/32

log syslog debugging
log facility syslog
```

4. In this directory, create the BGP configuration file using the following commands:

```
# emacs bgpd.conf
```

It should contain the appliance hostname, authentication details, response time, interfaces dedicated to BGP, access list and log file location like in the example below.

```
# more /data1/etc/quagga/bgpd.conf | grep v \!
hostname dns-anycast-1
password zebra
log syslog
router bgp 64500
bgp routerid 192.168.53.2
network 192.168.56.1/32
timers bgp 5 10
neighbor 192.168.53.1 remoteas 65000
neighbor 192.168.53.1 softreconfiguration inbound
neighbor 192.168.53.1 activate
```

To restart quagga

1. Open a shell session and connect to your appliance with *root* credentials.
2. Check the Quagga status using the following command:

```
/usr/local/etc/rc.d/quagga status
```

3. Restart Quagga using the following command:

```
/usr/local/etc/rc.d/quagga restart
```

Making Sure DNS Anycast Was Properly Configured for BGP

To make sure that DNS anycast is successfully implemented, on the router itself and on SOLIDserver, you can use a set of commands that ensure the IP address used during the configuration are part of the routes.

To display the BGP routes on Cisco routers

1. Connect to SOLIDserver via a shell session.
2. Run the following command:

```
router#show ip route
Codes: C connected, S static, I IGRP, R RIP, M mobile, B BGP
       D EIGRP, EX EIGRP external, O OSPF, IA OSPF inter area
       N1 OSPF NSSA external type 1, N2 OSPF NSSA external type 2
       E1 OSPF external type 1, E2 OSPF external type 2, E EGP
       i ISIS, su ISIS summary, L1 ISIS level1, L2 ISIS level2
       ia ISIS inter area, * candidate default, U peruser static route
       o ODR, P periodic downloaded static route

Gateway of last resort is not set

 192.168.56.0/32 is subnetted, 1 subnets
 B 192.168.56.1 [20/0] via 192.168.53.2, 00:02:57
 C 192.168.53.0/24 is directly connected, FastEthernet1/0
```

To display the BGP neighbors on SOLIDserver

1. Connect to SOLIDserver via a shell session.
2. Run the following command to connect to the zebra service:

```
# vtysh
Hello, this is Quagga (version 1.0.20160315).
Copyright 19962005 Kunihiro Ishiguro, et al.
dns1anycast#
```

3. Run the following command to display the routes.

```
dns1-anycast# show bgp neighbors
BGP neighbor is 192.168.53.1, remote AS 65000, local AS 64500, external link
  BGP version 4, remote router ID 192.168.53.1
  BGP state = Established, up for 00:01:41
  Last read 00:00:00, hold time is 10, keepalive interval is 3 seconds
  Neighbor capabilities:
    4 Byte AS: advertised
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    Graceful Restart Capabilty: advertised
  Message statistics:
    Inq depth is 0
    Outq depth is 0

```

	Sent	Rcvd
Opens:	2	2
Notifications:	1	0
Updates:	2	0
Keepalives:	50	43
Route Refresh:	0	0
Capability:	0	0
Total:	55	45

```

  Minimum time between advertisement runs is 30 seconds
  For address family: IPv4 Unicast
  Inbound soft reconfiguration allowed
  Community attribute sent to this neighbor(both)
  0 accepted prefixes
  Connections established 2; dropped 1
  Last reset 00:02:02, due to BGP Notification send
  Local host: 192.168.53.2, Local port: 27184
  Foreign host: 192.168.53.1, Foreign port: 179
  Nexthop: 192.168.53.2
  Nexthop global: ::
  Nexthop local: ::
  BGP connection: non shared network
  Estimated round trip time: 105 ms
  Read thread: on Write thread: off

```

To display the BGP summary on SOLIDserver

1. Connect to SOLIDserver via a shell session.
2. Run the following command to connect to the zebra service:

```
# vtysh
Hello, this is Quagga (version 1.0.20160315).
Copyright 19962005 Kunihiro Ishiguro, et al.
dns1anycast#
```

3. Run the following command to display the summary.

```
dns1-anycast# show bgp summary
IPv4 Unicast Summary:
-----
BGP router identifier 192.168.53.2, local AS Number 64500
RIB entries 1, using 112 bytes of memory
Peers 1, using 6960 bytes of memory
Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.53.1  4 65000    48     58     0    0    0 00:01:52      0
Total number of neighbors 1
```

Implementing Anycast Using IS-IS

IS-IS is a link-state routing protocol, operating by reliably flooding link state information throughout a network of routers. Successfully using anycast for IS-IS routing protocol on your network requires:

1. Meeting the [Prerequisites](#) to implement anycast.
2. [Configuring the Appliance for IS-IS Anycast](#) to make sure it uses the Quagga package that enables anycast.
3. [Configuring the Quagga Package for IS-IS Routing](#) to set the configuration that suits your needs.
4. [Making Sure DNS Anycast Was Properly Configured for IS-IS](#) once the configuration is complete.

Prerequisites

Anycast implementation prerequisites:

- Several servers in a pool must share 1 or anycast IPs.
- The servers must advertise their IP(s) to their neighboring routers.
- The routers exchange the routes information. That way if one server fails, the routers automatically recompile the routing tables to redirect the DNS clients.
- The 3 steps anycast configuration must be completed on all the appliances that manage a DNS server that you intend to include to the anycast routing scheme. This applies whether the servers are managed via a smart architecture or not.

With this type of topology, the anycast IP address is advertised from multiple locations and the router ends up choosing the best path to that IP address, according to the metric in use by the routing protocol. Once you finished the configuration detailed in the sections below, the DNS servers managed via SOLIDserver use anycast.

Configuring the Appliance for IS-IS Anycast

SOLIDserver contains a Quagga package that must be taken into account in the system configuration file to be used.

To successfully configure the package you must:

1. Meet the [Prerequisites](#).
2. Edit the rc.conf file to make sure it takes into account the package.
3. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*`.
4. Reboot the appliance. After the reboot, these files `running_conf.cf*` and `previous_conf.cf*` are created again and take into account the changes.
5. Configure the Quagga package. For more details, refer to the section [Configuring the Quagga Package for IS-IS Routing](#).

To configure the appliance for anycast DNS

1. Edit the system configuration file.
 - a. Open a shell session and connect to your appliance with `root` credentials.

- b. Open the file `/etc/rc.conf` to edit it.
- c. Enable Quagga and make sure the file is configured as follows:

```
quagga_daemons="zebra isisd"4
quagga_enable="YES"
```

- d. Add the following line to the file to specify the anycast dedicated IP address:

```
# Expected syntax: ifconfig_<interface-name>_alias0="inet <IP-address> netmask <netmask>"
# Example:
ifconfig_lo0_alias0="inet 192.168.99.8 netmask 255.255.255.0"
```

Save your changes.

2. Delete the files `/tmp/running_conf.cf*` and `/tmp/previous_conf.cf*` using the command:

```
rm /tmp/running_conf.cf* /tmp/previous_conf.cf*
```

3. Reboot the appliance using the command:

```
reboot
```

Now you need to configure the package following the section below.

Configuring the Quagga Package for IS-IS Routing

Once you configured the appliance to take into account the Quagga package, as detailed in the section [Configuring the Appliance for IS-IS Anycast](#), you can configure the package for IS-IS routing.

The package configuration implies:

1. Creating the Quagga and IS-IS dedicated configuration files.
2. Restarting Quagga.
3. Checking the logs.

To create the quagga dedicated configuration files

1. Open a shell session and connect to your appliance with *root* credentials.
2. Go to the directory `/data1/etc/quagga`.
3. In this directory, create the zebra configuration file using the following commands:

```
# emacs zebra.conf
```

It should contain the appliance hostname, administrator passwords, anycast IP address, anycast VIP(s) address and log file location like in the example below.

```
# more /data1/etc/quagga/zebra.conf | grep -v \!
hostname dns-anycast-1
password mypassword
enable password mypassword

log syslog debugging
log facility syslog
```

⁴To configure the appliance for both BGP and OSPF, you can specify the following: `quagga_daemons="zebra bgpd ospfd"`.

4. In this directory, create the IS-IS configuration file using the following commands:

```
# emacs isisd.conf
```

It should contain interfaces used for IS-IS communication and for anycast, an IS-IS address and log file location like in the example below:

```
# more /data1/etc/quagga/isisd.conf | grep v \!  
hostname dns-anycast-1  
password mypassword  
enable password mypassword  
  
interface em3  
 ip router isis ISIS_0  
  
interface lo0  
 ip router isis ISIS_0  
  
router isis ISIS_0  
 net 49.0001.1720.1600.2002.00  
 is-type level-1  
 metric-style wide  
  
log facility syslog
```

To restart quagga

1. Open a shell session and connect to your appliance with *root* credentials.
2. Check the Quagga status using the following command:

```
/usr/local/etc/rc.d/quagga status
```

3. Restart Quagga using the following command:

```
/usr/local/etc/rc.d/quagga restart
```

To check quagga log file

1. Open a shell session and connect to your appliance with *root* credentials.
2. In the file `/var/log/zebra.log` you can check the Quagga dedicated logs using the following command:

```
# tail -f /var/log/zebra.log
```

Making Sure DNS Anycast Was Properly Configured for IS-IS

To make sure that DNS anycast is successfully implemented, on the router itself and on SOLIDserver, you can use a set of commands that ensure the IP address used during the configuration are part of the routes.

To display the IS-IS routes and neighbors status on Cisco routers

1. Connect to SOLIDserver via a shell session.
2. Run the following command:

```
router> show ip route  
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, * - candidate default, U - per-user static route  
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```
+ - replicated route, % - next hop override
Gateway of last resort is not set
[...]
192.168.99.0/32 is subnetted, 1 subnets
i L1 192.168.99.8 [115/20] via 192.168.1.50, 00:00:33, FastEthernet0/0
[...]

router> show clns neighbors
System Id      Interface  SNPA                State Holdtime  Type Protocol
dns-anycast-1  Fa0/0     000c.29ef.d8bb     Up    28        L1   IS-IS
```

To display the IS-IS neighbors status on SOLIDserver

1. Connect to SOLIDserver via a shell session.
2. Run the following command to connect to the zebra service:

```
% vtysh
Hello, this is Quagga (version 1.1.1).
Copyright 1996-2005 Kunihiro Ishiguro, et al.
dns-anycast-1#
```

3. Run the following command to display info on IS-IS neighbor:

```
dns-anycast-1# show isis neighbor
Area ISIS_0:
System Id      Interface  L  State      Holdtime SNPA
router         em3       1  Up         8         ca01.025d.0000

dns-anycast-1# show isis hostname
Level System ID      Dynamic Hostname
1      1720.1600.1001 router
*      1720.1600.2002 dns-anycast-1
```

Integrating Cisco Umbrella

SOLIDserver embeds a DNSCrypt proxy which allows you to forward all the DNS queries it receives to the Cisco Umbrella Cloud⁵.

DNSCrypt is a protocol that authenticates communications between a DNS client and a DNS resolver. It prevents DNS spoofing. It uses cryptographic signatures to verify that responses originate from Cisco Umbrella and haven't been tampered with.

To successfully configure SOLIDserver for Cisco Umbrella:

1. Via Cisco Umbrella web interface, create a network device and retrieve its API key and secret strings.
2. Via SOLIDserver GUI, configure the IP address dedicated to Umbrella as the *only* forwarder for your local DNS appliance.
3. Via SOLIDserver CLI, configure and launch the proxy DNSCrypt.

Note that the DNSCrypt protocol uses the port 443, in TCP and UDP, which is usually reserved to HTTPS. It is possible that some equipments, such as firewalls, IDP or IPS detect a wrongful use of the port. Make sure these equipments are configured to allow this traffic.

⁵For more details on Cisco Umbrella services, refer to the proprietary website at <https://umbrella.cisco.com/>.

To configure the appliance for Cisco Umbrella

Only users of the group *admin* can perform this operation.

1. Retrieve your Cisco Umbrella parameters

- a. Connect to your Cisco Umbrella web interface using your credentials.
- b. In the left panel, click on **Admin > API Keys**. The page refreshes.
- c. Tick the box **Umbrella Network Devices**.
- d. Click on **CREATE**. The page refreshes and displays **Your Key** and **Your Secret** strings. Copy these values and keep them at hand as they disappear after you leave this page, you need them later on during the configuration.
- e. Tick the confirmation box and click on **CLOSE**.

2. Configure the DNS forwarder from the GUI

- a. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- b. At the end of the line of the server or smart architecture you intend to connect to Cisco Umbrella Cloud, click on **⚙**. The properties page opens.
- c. Open the panel **Forwarding** using **⌵**.
- d. Click on **EDIT**. The wizard **Forwarding configuration** opens.
- e. In the field **Add a forwarder**, type *127.0.1.53* which is the IP address dedicated to the proxy DNSCrypt.
- f. Click on **ADD** to move it to the list **Forwarders**.
- g. In the field **Forward mode**, select *Only*.
- h. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displayed the new settings.
- i. Repeat step *c* to *i* for each server or smart architecture on which you plan to deploy DNSCrypt.

3. Configure and launch the proxy DNSCrypt via CLI

- a. Open a shell session and connect to your appliance with *root* credentials.
- b. Retrieve your DNSCrypt parameters using the script *umbrella_setup*. You must specify the API key and API secret you copied earlier and the name of the device of your choice, as defined in your *Network Devices* list, as follows:

```
/usr/local/nessy2/script/umbrella_setup <Your-Cisco-Umbrella-API-Key>
<Your-Cisco-Umbrella-API-Secret> <Your-Network-Devices-Cisco-Umbrella-Device-Name>
```

The result should look as follows:

```
server 127.0.1.53 {
    edns-opendns yes;
    edns-opendns-orgid <your-Cisco-Umbrella-Organization-ID>;
    edns-opendns-deviceid "<your-Cisco-Umbrella-Device-ID>";
};
```

- c. Copy the lines returned.
- d. Edit the DNS *global* include file.

1. Open the `/data1/etc/namedb/global_include.conf`.
 2. Copy the lines of the DNSCrypt global parameters you retrieved.
 3. Save your changes.
- e. Edit the DNS *options* include file.
1. Open the `/data1/etc/namedb/options_include.conf`.
 2. Add the following line to the file, to specify DNSCrypt options parameters:

```
listen-on { !127.0.1.0/24; any ; };
```
 3. Save your changes.
- f. Edit the system configuration file.
1. Open the file `/etc/rc.conf`.
 2. Edit it to enable the proxy DNSCrypt as follows:

```
dnscrypt_proxy_enable="YES"
```
 3. Add the following line to the file to specify the IP address dedicated to the proxy DNSCrypt:

```
ifconfig_lo0_alias53="inet 127.0.1.53 netmask 255.0.0.0"
```
 4. Save your changes.
- g. Restart the network configuration using the command:

```
/etc/netstart
```
- h. Start the service `dnscrypt-proxy` using the command:

```
/usr/local/etc/rc.d/dnscrypt-proxy start
```
- i. Restart the service `ipmdns` using the command:

```
/usr/local/etc/rc.d/ipmdns.sh restart
```
- Now, every DNS query trafficking through the appliance is directly forwarded to the Cisco Umbrella Cloud for resolution using your organization ID and device ID, and therefore, your Umbrella policies.
- j. Repeat all the steps for each appliance you want to configure.

Chapter 37. Managing DNS Views

Within some DNS servers, you can create and administer views to serve one version of a zone to one set of clients and a different version of a zone to another set of clients. Views provide a different answer to the same DNS query, depending on the IP source of the query or the IP where the client packet is received. You can create multiple views of a given zone, with a different set of records in each of them. Same resource records can also exist in multiple zones in order to serve common records.

Browsing DNS Views

Within the DNS module, the view is the second level of the hierarchy. It allows you to manage zones, and therefore in extension, resource records. Keep in mind that also this level of the hierarchy is optional, **once you create views, all the zones have to be managed via a view whether all zones are managed through a unique view or several views.**

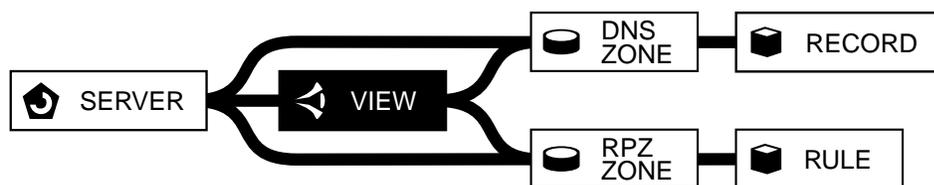


Figure 37.1. The view in the DNS hierarchy

Browsing the DNS Views Database

To display the list of DNS views

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. You can filter the list using the column search engines.

To display the DNS views of a specific server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the name of the server of your choice. The list of **All zones** of the server is displayed.
3. In the breadcrumb, click on **All views**. The list of views of the chosen server opens.

To display a view properties page

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **ⓘ**. The properties page opens.

Customizing the Display on the Page All Views

Users of the group *admin* can create customized column layouts. The button **⌵ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the DNS View Statuses

The column **Status** provides information regarding the views you manage.

Table 37.1. DNS view statuses

Status	Description
✔ OK	The view is operational.
🕒 Delayed create	The view is being created.
🗑️ Delayed delete	The view is being deleted.

You can also rely on the column **Multi-status** to monitor your view configuration with messages regarding the compatibility with Hybrid. For more details, refer to the section [Understanding the Column Multi-Status](#).

Adding DNS Views

You can add as many views as you need. Each view has a name and is configured with a list *match-clients* and/or *match-destinations*:

- **match clients** indicates which clients can access the view or not. It sets up a filter based on the IP address of the client requesting a specific resource. That way you can decide which IP address, or network, can access the zone(s) managed by the view.

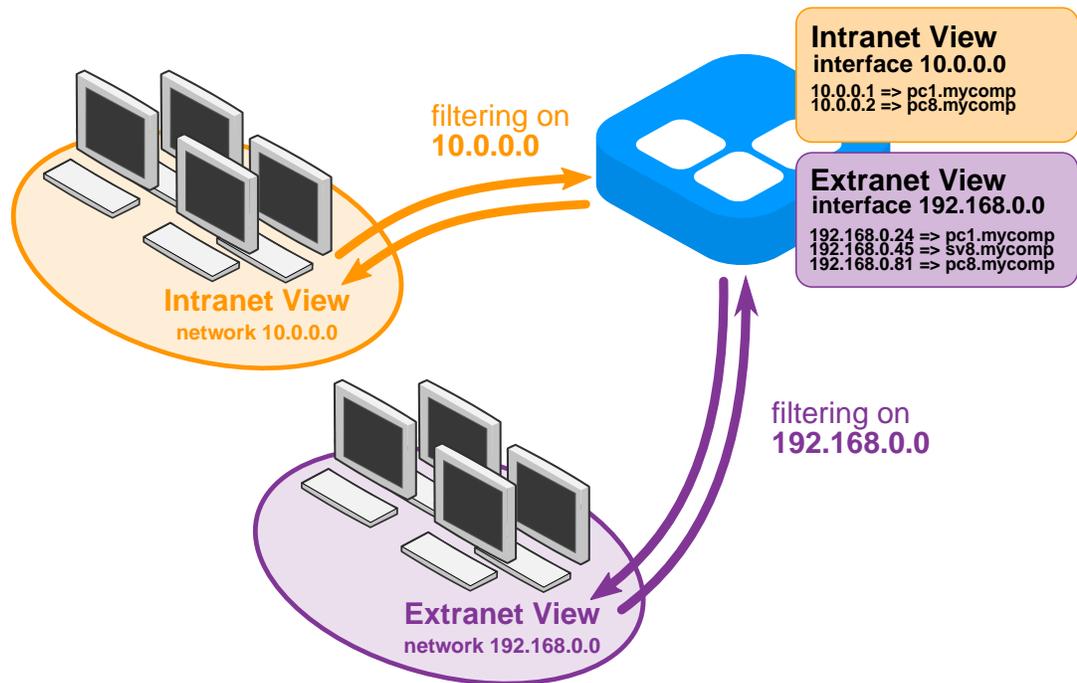


Figure 37.2. A DNS view configuration using match clients

- **match destinations** indicates where clients are directed. It sets up a filter based on the interface used by the client when querying the DNS server. Obviously, this criterion is only useful if you have several interfaces configured for one appliance.

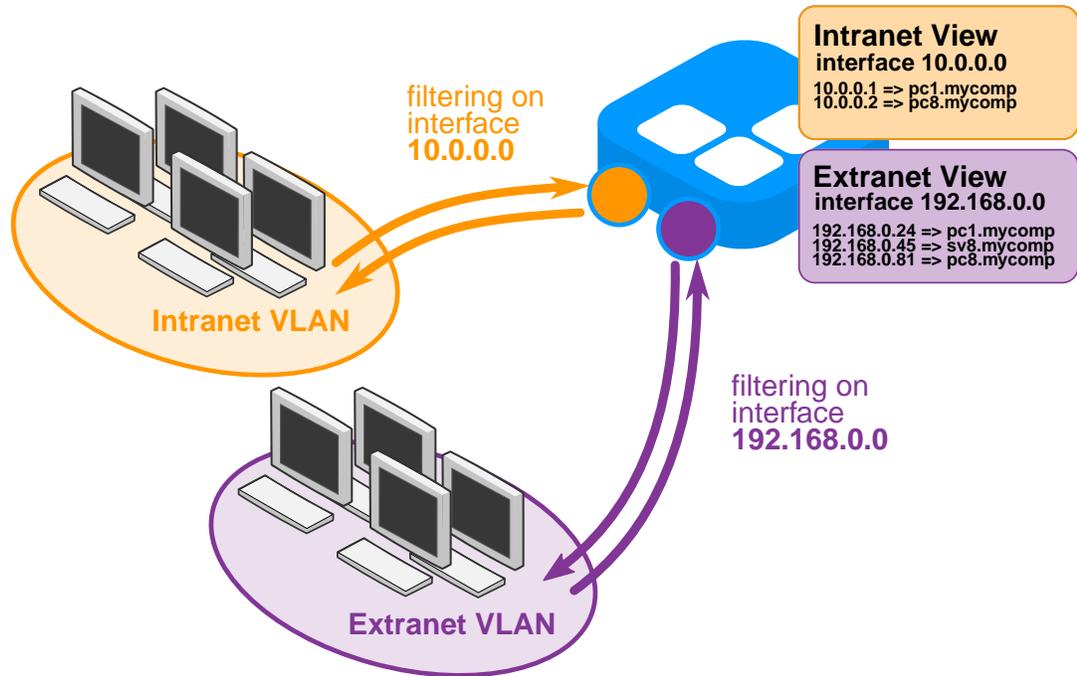


Figure 37.3. A DNS view configuration using match destinations

Keep in mind that **if you create views after creating zones, all the zones are moved to that view**. If you need several views, you have to create a another view and then move the zones in the view of your choice. For more details regarding zones migration, refer to the section [Copying or Moving DNS Zones](#). You cannot manage a set of zones through the views and others zones without the created views.

To add a DNS view

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. In the menu, select **+ Add > DNS view**. The **DNS server selection** wizard opens.
3. In the field **DNS server**, select the server on which you are adding a view.
4. Click on **NEXT**. The page **Add a DNS view** opens.
5. If you or your administrator created classes at the view level, in the list **DHCP view class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **DNS view name**, type in an explicit name. This name cannot contain special characters. It can contain letters and numbers, for instance *external*, *internal1* and *internal2* are correct view names.
7. Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below.

Table 37.2. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

If you want to create a view and configure it later, click on **NEXT** until to get to the last page of the wizard and click on **OK** to commit the creation. For more details regarding the default configuration settings of a view, refer to the last step of this procedure.

8. Click on **NEXT**. The page **Match clients** opens.
9. You can configure the list **ACL values**:
 - a. Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 37.3. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

- b. Once a restriction/permission is configured as needed, click on **ADD**. The value is moved to the list **ACL values**. In the list, denied hosts are preceded by an exclamation mark (!).

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove the entry from the list. If you made changes that you do not want to save, click on **CANCEL**.
 - c. Repeat these actions for as many restrictions and permissions as needed.
 - d. Make sure the order of the values in the list **ACL values** suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. If not, select a value and use the buttons **▲** and **▼** to organize the list.

10. Click on **NEXT**. The page **Match destinations** opens.
11. You can configure the list **ACL values** as detailed in step 9.

- Click on **NEXT**. The page **DNS views order** opens.
- In the field **DNS views order**, the view you are creating is listed. It is always listed at the bottom of the list.

Once you created more than one view, you can order the list using  and . Within a server, each view *match clients* and *match destinations* is reviewed following the order set in this field. For more details, refer to the section [Editing the Order of the Views](#).

- Click on **OK** to complete the operation. The report opens and closes. Note that by default, if you do not configure any ACL list for the view:
 - The **match clients** is configured with a key named *key <viewname>*.
 - The **match destinations** is configured with a key *key <viewname>* and the ACL *any*. If you do not edit or delete this ACL, it grants access to anyone.

The page refreshes, the configuration is visible in the columns **Match clients** and **Match destinations**.

Adding a view automatically edits the match clients and match destinations of the existing view(s), with a key *! key <newviewname>*. This ensures that the views deny access to each other and manage separate zones and RRs.

Every time you add a new view, the **Status** of all the views changes from  *OK* to  *Delayed create* while their *match clients* is being edited. Once it is done, they all change back to *OK*.

Editing DNS Views

Once a view is added, you can edit it. Before editing a view, keep in mind that:

- You cannot edit the name of a view.**
- The list *match-clients* allows to grant or restrict access to the zones managed by the view based on the source IP address of the incoming DNS requests.
- The list *match-destinations* allows to set the destination address¹ of the incoming DNS requests based on the interface used for the request. If you did not set several interfaces, it is useless to configure the match-destination.
- The lists *match-clients* and *match-destinations* are access control lists in essence. Therefore, the order of the elements listed in both lists is important as each restriction or permission is reviewed following the order you set in the list.

To edit a DNS view

- In the sidebar, go to  **DNS > Views**. The page **All views** opens.
- Put your mouse over the **Name** of the view you want to edit. The contextual menu appears.
- Click on . The wizard **Edit a DNS view** opens.
- If you or your administrator created classes at the view level, in the list **DHCP view class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

¹The destination IP address is actually a DNS server interface.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the advanced properties or the value of the drop-down list **Advanced properties** if need be. For more details, refer to the section [Configuring DNS Advanced Properties](#).
6. Click on **NEXT**. The page **Match clients** opens.
7. Edit the list **ACL Values** according to your needs. For more details regarding the available settings, refer to the procedure [Configure the list ACL values](#).
 - To add an entry, fill in the fields as detailed in the procedure [To add a DNS view](#).
 - To update an entry, select it in the list. Its information is loaded in the fields above. Make your changes and click on **UPDATE**.
 - To discard the latest modifications, click on **CANCEL**.
 - To reorganize the order of the entries in the list, select them one by one in the list and click on or until the order suits your needs. Each restriction or permission is reviewed following the order you set in the list.
 - To delete an entry, select it in the list and click on **DELETE**.
8. Click on **NEXT**. The page **Match destinations** opens.
9. Edit the list **ACL Values** as detailed in step 7.
10. Click on **NEXT**. The page **DNS views order** opens.
11. In the field **DNS views order**, order the views according to your needs using and . For more details, refer to the section [Editing the Order of the Views](#) below.
12. Click on **OK** to complete the operation. The report opens and closes. The page refreshes, the changes are visible in the columns **Match clients** and **Match destinations**.

Editing the Order of the Views

Once you created several views on a server, you can order them. The Order set is displayed in the column **Order**. If you only have one view on a server, its value is 0.

Ordering views on a server allows to specify in which order the match client and match destination configurations of each view (ACL, networks, etc.) are reviewed. This, in turn, impacts the DNS client queries responses. The order of the views you set is followed strictly: once a match is found, the rest of the restrictions and permissions are ignored. The first view reviewed is 0, the second on is 1, and so forth. This order is saved in the DNS server configuration.

To edit the order of the views

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. Click on the name of the smart server of your choice. Only the views of the selected server are displayed.
3. Right-click over the **Name** of the view of your choice. The contextual menu opens.
4. Click on **🔗**. The wizard **Add a DNS view** opens.
5. The field **DNS view name** is gray to indicate that it cannot be edited.

6. In the drop-down list **Advanced properties**, edit the value and corresponding fields if need be. For more details, refer to the section [Configuring DNS Advanced Properties](#).
7. Click on . The page **Match clients** opens.
8. Click on . The page **Match destinations** opens.
9. Click on . The page **DNS views order** opens.
10. In the field **DNS views order**, order the views according to your needs using and .
11. Click on to complete the operation. The report opens and closes. The page refreshes. The new order set is visible in the column **Order**.

Deleting DNS Views

At any moment you can delete a view. Before deleting a view, keep in mind that:

- The views must be **deleted one by one**.
- Deleting one view deletes the zone(s) it manages, as well as all the RRs the zone(s) manage on the physical server. So if you want to delete a view but not the zones it contains, migrate the zones to a different view before deleting it. For more details regarding zones migration, refer to the section [Copying or Moving DNS Zones](#).
- Deleting a view, removes it from the DNS views order list: the list is updated. This order is also updated in the DNS configuration.
- If you only have one view, deleting it does not delete the zone(s) it manages but only the container itself: the view is therefore no longer listed on the server All views page.
- If zones inherited class parameters from the deleted view, their value and the source of their value remain the same: the *Inheritance property* of each class parameter is forced to *Inherit* or *Set* to match the configuration of the deleted view.

If you want get rid of all the views and manage zones via the DNS server itself, refer to the section [Going Back to Managing Zones Without Views](#).

To delete a view

1. In the sidebar, go to  **DNS > Views**. The page **All views** opens.
2. Filter the list if need be.
3. Tick the view you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The view is Deleted before it is no longer listed. In the meantime, the zones and RRs it managed are deleted as well if you had several created views.

Defining a DNS View as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a view as one of the resources of a specific group allows the users of that group to manage the view in question as long as they have the corresponding rights granted.

Granting access to a view as a resource also makes every item it contains available. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Going Back to Managing Zones Without Views

At any time you might want to stop managing your zones with views. Considering that the way you delete views has an impact on the database and different behaviors you need to be careful. First, keep in mind that no matter how many views you created, the last view listed on the page *All views* of a specific server can be deleted on its own: it does not delete the zones it manages.

With that in mind, we recommend that you follow the steps below to successfully get rid of the views when you no longer need them.

To successfully remove all views

1. Choose the view that should be deleted last.
2. Migrate all the zones you want to keep in that view. For more details regarding zones migration, refer to the section [Copying or Moving DNS Zones](#).
3. One by one, tick and delete the unwanted views. For more details, refer to the procedure [To delete a view](#).
4. Once the only remaining view is the one that holds all the zones you want to work with, tick it and delete it. The zones and RRs it contains are kept and still listed in the pages *All zones* and *All RRs* of the server. Now you can manage them through the server directly.

Chapter 38. Configuring DNS Views

Like servers, views can be configured individually to set a series of behaviors for the zones they contain.

Any configuration set at view level overwrites what was set at server level (whether physical or smart).

Configuring DNS Forwarding at View Level

You can set a forwarding configuration on a view from its properties page. Note that:

- On a view belonging to a server not managed via a smart architecture, the specific forwarding configuration only applies to the view on said server.
- On a view belonging to a server managed via a smart architecture, the specific forwarding configuration applies to the view on all the servers managed by the smart.
- The forwarding configuration set on a smart architecture view is automatically inherited by the zones it manages. You can override the configuration directly on the physical server view.
- Any configuration set at view level overrides the server level configuration. All the zones managed by the view inherit the new settings.
- **Any configuration set at zone level overrides the view level configuration.** For more details, refer to the section [Configuring DNS Forwarding at Zone Level](#).

To configure forwarding at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Forwarding** using **⌵**.
4. Click on **EDIT**. The wizard **Edit a DNS view** opens.
5. Click on **NEXT** until the page **Forwarding configuration** appears.
6. In the drop down-list **Forward mode**, select the mode of your choice according to the table below.

Table 38.1. Forward mode options at view level

Option	Description
Default	The view uses the forward configuration set at the server level.
None	The forwarding is disabled.
First	The server sends the queries to the forwarder(s) configured for the view and, if it does not receive any answer, attempts to find an answer itself.
Only	The server only forwards queries to the forwarder(s) configured for the view.

7. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

8. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displayed the new settings. In the panel **Forwarding**, your configuration is displayed.

You can set a specific forwarding configuration for a view belonging to a physical server managed via a smart architecture. This new configuration is inherited by the zones and records of the view. Keep in mind that:

- When a forward mode is set on a smart architecture view, you cannot set the forward mode to *None* on a view belonging to a physical server managed via a smart architecture. You can only set a different forward mode.
- Any configuration set at zone level overrides the view level configuration.

To configure a specific forward mode on a physical server view

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. Make sure the servers managed by your smart architectures are displayed. If not, on the right-end side of the menu, click on .
3. At the end of the line of the view of the physical server of your choice, click on . The properties page opens.
4. Open the panel **Forwarding** using .
5. Click on **EDIT**. The wizard **Edit a DNS view** opens.
6. Click on **NEXT** until the page **Forwarding configuration** appears.
7. Tick the box **Overwrite the smart settings**. The page refreshes and displays additional fields.
8. In the drop down-list **Forward mode**, you can select *First* or *Only*. For more details refer to the table [Forward mode options](#).
9. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

10. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displayed the new settings. In the panel **Forwarding**, the *Forward* value is preceded by the message *Smart configuration is overwritten*.

To revert the specific configuration and inherit it again, edit the *Forwarding* to untick the box *Overwrite the smart settings*.

Configuring DNS Notify Messages at View Level

Configuring the Notify at view level allows to set the changes notification for all the master zones it contains. Once the notification is sent to slave zones, the administrator decides if a zone transfer is relevant. For more details, refer to the section [Limiting Zone Transfer at View Level](#).

Within SOLIDserver, the notification configuration is done from the panel *Notify* of the properties page. This panel displays:

- The notification type configured for the view,
- The slave zones that receive the notify messages through their managing view (Also notify),
- The allow-notify directive of the view slave zones. For instance, you can allow all the servers of a network to notify the slave zones of your server or only a few.

Note that there is an implicit allow-notify directive set when you add a slave zone: when you set the Master IP address of the slave zone you are allowing the master zones of this server to send notify messages to your slave zone.

Keep in mind that **any configuration set at view level overrides the server level configuration and any configuration set at zone level overrides the view level configuration.**

To configure notify messages at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **ⓘ**. The properties page opens.
3. Open the panel **Notify** using **Ⓜ** and click on **EDIT**. The wizard opens.
4. If you or your administrator created classes, the **DNS view class** list is visible. Select a class or *None* and click on **NEXT**. The next page of the wizard opens.
5. Do not edit the advanced properties configuration and click on **NEXT**. The page **Notify** opens.
6. In the drop-down list **Notify**, configure the view notification behavior following the table below.

Table 38.2. DNS view notify types

Field	Description
No	With this option, no notify message is sent when changes are performed in the master zones.
Yes	With this option, the notify messages is sent to the target of the NS records of the master zone. It is also sent to the IP address(es) specified in the field <i>IP address</i> below.
Explicit	With this option, the notify messages is only sent to the IP address(es) specified in the field <i>IP address</i> below.

7. If you selected *Yes* or *Explicit*, you can set the IP address and port of the server(s) which slave zones should receive the messages:
 - a. In the field **IP address**, type in the IP address of another server. The notify message is sent if you chose the notify type *Yes* or *Explicit*.
 - b. In the field **Port**, you can specify the port number that should receive the notify messages on the server you specified in the previous field.

- c. Click on **ADD**. The IP address and port number are displayed in the list **Also notify** as follows: *<ip-address> port: <port-number>*. You can repeat these actions for as many servers as needed.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove the entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- 8. Click on **NEXT**. The page **Allow notify** opens. It allows to specify if the view slave zones can receive master zones notification messages.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 38.3. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- 9. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again. Your configurations are displayed in the panel **Notify**.

Configuring DNS Recursion at View Level

The recursion settings at server level are inherited by the views. However, you can change these settings at view level to customize the recursion configuration on the network: the changes operated on view are inherited by the zones managed through the view.

Enabling and Disabling the Recursion on a View

The recursion statement essentially controls caching behavior in the view and the zones it manages.

From the view properties page, you can edit its recursive behavior through the panel *Recursion*. By default, its content is inherited from the server.

Keep in mind that **any configuration set at view level overrides the server level configuration.**

To enable the DNS recursion on a view

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb, click on **All views**. The page **All views** of the physical server opens.
3. At the end of the line of the view of your choice, click on **⌵**. The properties page opens.
4. Open the panel **Recursion** using **⌵**.
5. Click on **NEXT** until you get to the page **Recursion configuration**.
6. Open the panel **Recursion** using **⌵**. If the **Recursion** is set to *no*, click on **EDIT**. The wizard opens.
7. Click on **NEXT** until you open the page **Recursion configuration**.
8. In the drop-down list **Recursion**, select *yes*.
9. Click on **NEXT**. The page **Allow recursion** opens. For more details regarding the recursion configuration, refer to the section [Limiting the Recursion at View Level](#) below.
10. Click on **OK** to complete the operation.

To disable the DNS recursion on a view

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb, click on **All views**. The page **All views** of the physical server opens.
3. At the end of the line of the view of your choice, click on **⌵**. The properties page opens.
4. Open the panel **Recursion** using **⌵**. If the **Recursion** is set to *yes*, click on **EDIT**. The wizard opens.
5. Click on **NEXT** until you open the page **Recursion configuration**.
6. In the drop-down list **Recursion**, select *no*.
7. Click on **OK** to complete the operation. The report opens and closes. The page refreshes, in the panel the recursion is disabled.

Limiting the Recursion at View Level

By default, the view inherits the server recursion settings (permissions and restrictions).

Once the recursion is restricted at view level, it overrides the server recursion configuration and applies to the zones it contains.

To set an allow-recursion match list at view level

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb, click on **All views**. The page **All views** of the physical server opens.
3. At the end of the line of the view of your choice, click on **⌵**. The properties page opens.
4. Open the panel **Recursion** using **⌵** and click on **EDIT**. The wizard opens.
5. Click on **NEXT** until you get to the page **Recursion configuration**.

- Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 38.4. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Restricting DNS Queries at View Level

The DNS queries can be restricted through the allow-query and allow-query-cache options. They both set an ACL list for IP addresses and/or network addresses, so keep in mind that **the order of the elements listed in the field ACL values is important** as each restriction or permission is reviewed following the order you set in the list.

Allow Query

You can specify which hosts are allowed to issue DNS queries.

The allow-query configuration set at view level overrides the allow query defined at server level. Once the statement is set for a view, it applies to all the zones it contains.

Keep in mind that **any configuration set at view level overrides the server level configuration and any configuration set at zone level overrides the view level configuration**.

To set an allow query match list at view level

- In the sidebar, go to **DNS > Views**. The page **All views** opens.
- At the end of the line of the view of your choice, click on **⊞**. The properties page opens.
- Open the panel **Access control** using **⊞**. This panel displays different options: **Allow-query**, **Allow query cache** and **Allow-transfer**.

4. Click on **[EDIT]**. The wizard **Add a DNS view** opens.
5. In the field **DNS view name**, the view name is displayed in gray to indicate you cannot edit it.
6. In the drop-down list **Advanced properties**, you can select *Default* or *All* and edit the configuration if need be. For more details, refer to the chapter [Managing Advanced Properties](#).
7. Click on **[NEXT]**. The page **Allow-query** opens.
8. Set up the allow-query match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 38.5. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **[ADD]**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **⬅** and **➡**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **[UPDATE]** or click on **[DELETE]** to remove an entry from the list. If you made changes that you do not want to save, click on **[CANCEL]**.

9. Click on **[NEXT]** twice to skip the page **Allow-query-cache** and open the page **Allow-transfer**.
10. Click on **[OK]** to complete the operation. The report opens and closes. The properties page is visible again, your configuration is listed in the list **Allow-query** of the panel **Access control**.

Allow Query Cache

You can specify which hosts are allowed to issue DNS queries on the local view cache.

The allow query cache configuration set at view level overrides the allow query cache defined at the server level. Once the statement is set for a view, it applies to all the zones it contains.

Allow-query-cache statement particularities

The allow-query-cache is independent from the allow-query statement but closely linked to the allow-recursion statement.

If the recursion is set to no, the cache cannot be queried, so it is useless to set an allow-query-cache match list.

If the recursion is set to yes and the allow-recursion statement is not defined, by default the localhost and localnets are permitted to query the server cache.

If the recursion is set to yes and the allow-recursion statement is defined with a specific match list, the local cache access is granted to all the entries of the allow-recursion match list.

The match list defined controls recursive behavior as recursive queries would be useless without access to the local view cache. Typically, if a host is in the allow-recursion match list, it could access the view the first time and get query result. However, if it is not part of the allow-query-cache match list then it would not be able to make the same query a second time as it would be saved on the cache to which it does not have access. On the contrary, if a host is in the allow-query-cache match list but not in the allow-recursion match list, it would only get results for queries already sent by another host with the proper access rights. Hence the need to configure carefully both these statements to avoid conflicts and absurd access configurations.

To set an allow query cache match list at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **ⓘ**. The properties page opens.
3. Open the panel **Access control** using **⊞**. This panel displays different options: **Allow-query**, **Allow query cache** and **Allow-transfer**.
4. Click on **[EDIT]**. The wizard **Add a DNS view** opens.
5. In the field **DNS view name**, the view name is displayed in gray to indicate you cannot edit it.
6. In the drop-down list **Advanced properties**, you can select *Default* or *All* and edit the configuration if need be. For more details, refer to the chapter [Managing Advanced Properties](#).
7. Click on **[NEXT]**. The page **Allow-query** opens.
8. Click on **[NEXT]**. The page **Allow query cache** opens.
9. Set up the allow-query-cache match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 38.6. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **[ADD]**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed

is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons  and .

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

10. Click on **NEXT**. The page **Allow-transfer** opens.
11. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again, your configuration is listed in the list **Allow-query-cache** of the panel **Access control**.

Limiting Zone Transfer at View Level

DNS zone transfer is a type of DNS transaction employed to replicate and synchronize all copies of the zone used at each server configured to host the zone. SOLIDserver denies zone transfers by default to all DNS server but supports the allow-transfer property at view level to allow you to specify which hosts, networks, or TSIG keys are granted or denied the permission to do transfers for all the zones of the view.

The allow-transfer option configuration basically creates an ACL dedicated to controlling transfers so keep in mind that **the order of the elements listed in the field ACL values is important** as each restriction or permission is reviewed following the order you set in the list.

The allow-transfer property set at view level overrides the allow query defined at server level. Once the statement is configured at view level, it applies to all the zones it contains.

Keep in mind that **any configuration set at view level overrides the server level configuration and any configuration set at zone level overrides the view level configuration**.

To set an allow transfer match list at view level

1. In the sidebar, go to  **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Access control** using . This panel displays different options: Allow-query, Allow query cache and Allow-transfer.
4. Click on **EDIT**. The wizard **Add a DNS view** opens.
5. In the field **DNS view name**, the view name is displayed in gray to indicate you cannot edit it.
6. In the drop-down list **Advanced properties**, you can select *Default* or *All* and edit the configuration if need be. For more details, refer to the chapter [Managing Advanced Properties](#).
7. Click on **NEXT**. The page **Allow-query** opens.
8. Click on **NEXT**. The page **Allow query cache** opens.
9. Click on **NEXT**. The page **Allow-transfer** opens.
10. Set up the allow-transfer match list.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 38.7. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost and localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

11. Click on **NEXT**. The page **Allow-transfer** opens.
12. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again, your configuration is listed in the list **Allow-transfer** of the panel **Access control**.

Configuring Client Resolver Cache Options at View Level

From the properties page of a view belonging to a smart architecture managing EfficientIP DNS servers using the SSL protocol, you can edit the *lame-ttl* and *max-cache-client* options.

These options set at view level override the server level configuration. Once they are configured at view level, they apply to all the zones it contains.

For more details regarding these two options, refer to the section [Configuring Client Resolver Cache Options at Server Level](#).

To set the *lame-ttl* option at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Options** using **⌵** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Lame-ttl**, type in the value of your choice. This value is in seconds can be set between 30 and 1800. The default value is 600, the maximum value is 1800 seconds.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

To set the *max-cache-size* option at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.

2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Options** using  and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Max-cache-size**, type in the value of your choice to set the cache memory size. This value is in bytes. The default value is *100m*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

Configuring EDNS Options at View Level

From the properties page of a view belonging to a smart architecture managing EfficientIP DNS servers using the SSL protocol, you can edit the *edns-udp-size* and *max-udp-size* options.

These options set at view level override the server level configuration. Once they are configured at view level, they apply to all the zones it contains.

For more details regarding these options, refer to the section [Configuring EDNS Options at Server Level](#).

To set the edns-udp-size option at view level

1. In the sidebar, go to  **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Options** using  and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Edns-udp-size**, type in the size of received packets of your choice. This value is in bytes, and must be set between 512 and 4096. The default value is *4096*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

To add a DNS key at view level

1. In the sidebar, go to  **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Options** using  and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Max-udp-size**, type in the maximum size of the packets you send. This value is in bytes and must be set between 512 and 4096. The default value is *4096*.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the value you defined.

Configuring a Sortlist at View Level

From the properties page of a view belonging to a smart architecture managing EfficientIP DNS servers using the SSL protocol, you can edit the *sortlist* statement can be edited at view level. Like any other configuration option, the settings defined at server level are edited by the view. Editing them at view level overwrites the server level configuration and applies to the zones managed by the view.

For more details regarding the sortlist statement, refer to the section [Configuring a Sortlist at Server Level](#).

Keep in mind that **any configuration set at view level overrides the server level configuration.**

To define a sortlist statement at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **⊞**. The properties page opens.
3. Open the panel **Options** using **⊞** and click on **EDIT**. The wizard **Options configuration** opens.
4. In the field **Client address**, type in the client IP address/subnet. It must be composed of an IPv4 address containing 1 to 4 bytes followed by the prefix: `<IP address>/<prefix>`.
5. In the field **Sort address**, type in a list of IP addresses or subnets followed by a semi-colon. These addresses correspond to the value of an A record of the RRset for which you create the sortlist. The statement respects the order in which you typed in the addresses. The value must respect the format `<IP address>/<prefix>;` even if you only type in one sort address.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays your sortlist as follows: `{<client_address_field_value>;<first_sort_address>;<second_sort_address>;<etc>;}`. There is one sortlist per client address defined.

Configuring DNS Sources at View Level

The Sources configuration is only available for views managed through an EfficientIP DNS physical server using the SSL protocol.

Configuring DNS sources allows to set physical interfaces at view level to be systematically used for all notify operations and zone transfer. DNS sources configuration can be inherited from the server. If set at view level, it is inherited by the zones. The inheritance details are visible on both the views and zones properties page.

Keep in mind that **any configuration set at view level overrides the server level configuration and any configuration set at zone level overrides the view level configuration.**

From the **Sources** and **Sources V6** panels, through their IP address, you can configure physical interfaces to be used for the view transfer and notify options. When editing these panels, you can define the following statements:

query-source

This statement allows to define the IPv4 address and/or port used as the source of the server or view outgoing queries. By default, BIND uses any server or view interface IP address and a random port for outgoing queries.

Using a fixed port number allows to control UDP operations but can be extremely dangerous: it can lead to cache poisoning if used with any caching DNS server definition as any attacker would need to guess the transaction ID to get both the specified interface IP address and port number. This statement is displayed on servers and views properties page.

query-source-v6

This statement allows to define the IPv6 address and/or port used as the source of the server or view outgoing queries. By default, BIND uses any server or view interface IP address and a random port for outgoing queries.

Using a fixed port number allows to control UDP operations but can be extremely dangerous: it can lead to cache poisoning if used with any caching DNS server definition as any attacker

would need to guess the transaction ID to get both the specified interface IP address and port number. This statement is displayed on servers and views properties page.

transfer-source

This statement allows to determine the IPv4 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

transfer-source-v6

This statement allows to determine the IPv6 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

use-alt-transfer-source

This statement allows to set the use of an alternate interface IP address for the transfer if the *transfer-source* or the *transfer-source-v6* were to fail. This statement configuration is displayed on the physical server, view and slave zones properties page.

This statement definition is only configurable from the panel **Sources** but applies to interfaces whether they were identified through an IPv4 or an IPv6 address.

Its default value is *no* if the server contains views and *yes* if the server does not contain any view.

alt-transfer-source

This statement allows to determine the alternate IPv4 address of the interface used to execute the zones transfer on the server if the *transfer-source* fails and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

alt-transfer-source-v6

This statement allows to determine the alternate IPv6 address of the interface used to execute the zones transfer on the server if the *transfer-source* fails and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

notify-source

This statement allows to define the IPv4 address of the physical interface used for all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

notify-source-v6

This statement allows to define the IPv6 address of the physical interface used all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

To set IPv4 DNS sources at view level

1. In the sidebar, go to  **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Sources** using  and click on . The wizard **Configuration: Sources** opens.

4. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise, all the transfer operations would fail.
 - a. In the field **Query-source address**, type in the IPv4 address of the interface used for outgoing queries.
 - b. In the field **Query-source port**, you can type in the port number used for outgoing queries. Keep in mind that specifying a port number can lead to cache poisoning if DNS server definitions are not set properly.
 - c. In the field **Transfer-source address**, type in the IPv4 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance.
 - d. In the field **Transfer-source port**, you can specify which port on the interface should be used.
 - e. In the drop-down list **Use-alt-transfer-source**, set the use of an alternate interface if need be.

Table 38.8. Use-alt-transfer-source parameters

Parameter	Description
none	This is the default value of the <i>use-alt-transfer-source</i> statement. If your server contains views, it corresponds to <i>no</i> . If your server does not contain any view, it corresponds to <i>yes</i> .
no	This value disables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails. Go to step 5 to set the notify-source statements related fields.
yes	This value enables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails. In this case, you need to set the alternate interface IP address (and port if you want) through the <i>alt-transfer-source</i> and <i>alt-transfer-source-v6</i> statements in the following steps.

The statement *use-alt-transfer-source* applies to the alternate interfaces declared through IPv4 address (Alt-transfer-source address) and IPv6 address (Alt-transfer-source address-v6).

- f. If you enabled the use of an alternate interface, in the field **Alt-transfer-source address**, type in the IPv4 address of the alternate interface. It must also be configured on the appliance.
 - g. If you enabled the use of an alternate interface, in the field **Alt-transfer-source port**, you can specify which port on the interface should be used.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the notify operations would fail.
 - a. In the field **Notify-source address**, type in the IPv4 address to be used for outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source port**, you can specify which port on the interface should be used.
6. Click on to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

To set IPv6 DNS sources at view level

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.

2. At the end of the line of the view of your choice, click on . The properties page opens.
3. Open the panel **Sources** using  and click on . The wizard **Configuration: Sources** opens.
4. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the transfer operations would fail.
 - a. In the field **Query-source-v6 address**, type in the IPv6 address of the interface used for outgoing queries.
 - b. In the field **Query-source-v6 port**, you can type in the port number used for outgoing queries. Keep in mind that specifying a port number can lead to cache poisoning if DNS server definitions are not set properly.
 - c. In the field **Transfer-source-v6 address**, type in the IPv4 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance. If you defined the *use-alt-transfer-source* statement in the panel **Sources**, it applies to the alternate interfaces declared in IPv4 (Alt-transfer-source address) and IPv6 (Alt-transfer-source address-v6).
 - d. In the field **Transfer-source-v6 port**, you can specify which port on the interface should be used.
 - e. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 address**, type in the IPv6 address of the alternate interface. It must also be configured on the appliance.
 - f. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 port**, you can specify which port on the interface should be used.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the notify operations would fail.
 - a. In the field **Notify-source-v6 address**, type in the IPv6 address to be used for the outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source-v6 port**, you can specify which port on the interface should be used.
6. Click on  to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

Chapter 39. Managing DNS Zones

When deploying a name server, it is important to understand the difference between a zone and a domain. A zone is a delegated point within a DNS structure, and is made up of adjoining elements of the domain structure, which are governed by a name server.

SOLIDserver allows you to create and manage 6 types of zones: *Master*, *Slave*, *Forward*, *Stub*, *Hint* and *Delegation-Only*. Each type of zone provides specific options that you can set when creating or editing the zones.

If you want to create RPZ zones, refer to the chapter [DNS Firewall \(RPZ\)](#).

Browsing DNS Zones

As far as the DNS hierarchy is concerned, the zone is the third level. It is compulsory to create a zone to manage resource records.

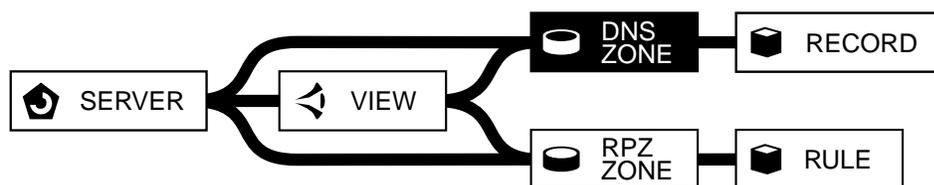


Figure 39.1. The zone in the DNS hierarchy

Browsing the DNS Zones Database

To display the list of zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. To display the list of zones of a specific server, in the column **Server**, click on the name of the server of your choice. The page refreshes.

To display a zone properties page

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **[icon]**. The properties page opens.

On the DNS zone properties page of a physical server you can find the following information in separate panels:

- **Main properties:** sums up the main information regarding the zone. In the case of our zone: the zone name, type, resolution, server, view, IPAM space it is linked to, responsible user email address, refresh frequency, lifespan, etc. In other words, everything there is to know about the zone apart from specific configurations that are all displayed in dedicated panels.
- **Advanced properties:** displays the advanced properties set at zone level. You cannot edit them from this panel, to make any changes use the button **EDIT** in the panel **Main properties** and change them on the dedicated page.

- **Name servers:** displays the server(s) that have authority over the zone or over the domain sub-zone(s).
- **Forwarding:** displays the servers toward which are redirected the DNS queries for that zone.
- **Groups access:** displays the groups that have the zone listed as a resource and the rights its users have over it.
- **Ticket:** displays the users that issued a ticket through the Workflow module to modify or delete the zone. This panel cannot be edited from the properties page.
- **Notify:** displays the IP addresses of the servers that should be notified of any change made on the master zone. These servers contain slave zone(s) named after the current master zone.
- **Access control:** displays the allow-query, allow-transfer and allow-update access permissions and restrictions to query the master zone, transfer the zone data or update the zone.
- **Sources:** displays the IPv4 interface(s) used to send the zone notifications.
- **Sources V6:** displays the IPv6 interface(s) used to send the zone notifications.
- **State log:** displays the server status evolution log; OK or KO (i.e. Timeout) and at what time it changed status. This panel cannot be edited, it simply provides information.
- **Audit:** displays all of the latest changes performed on the zone by the user logged in. If they belong to a group with access to the changes from all users, the panel displays all the operations ever performed. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).

Customizing the Display on the Page All Zones

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the DNS Zone Statuses

The column **Status** provides information regarding the zones you manage.

Table 39.1. DNS zone statuses

Status	Description
 <i>OK</i>	The zone is operational.
 <i>Busy</i>	The zone is synchronizing.
 <i>Delayed create</i>	The zone creation is delayed due to a server load or a server unavailability. The creation is automatically pushed when the server is available.
 <i>Delayed delete</i>	The zone deletion is delayed due to a server load or a server unavailability. The deletion is automatically pushed when the server is available.
 <i>Timeout</i>	The zone is not available.
 <i>Unknown</i>	The zone is not synchronized yet.
 <i>Not authoritative</i>	The zone configuration is incorrect: in the SOA another server was set as authoritative.
 <i>Refused</i>	The DNS server refuses the transfer between the current zone and the management platform, check the parameter <i>allow-transfer</i> of the zone or server properties page.
 <i>No RR</i>	There is no RR to transfer for the zone. That status can be displayed for a <i>forward</i> zone.
 <i>Unmanaged</i>	The zone is not available due to a disabling operation.

Moreover, the column **Multi-status** provides you with emergency, warning, critical, error or informational messages regarding the compatibility with Hybrid. For more details, refer to the section [Understanding the Column Multi-Status](#).

Adding DNS Zones

You can add as many zones as you want on the page *All zones*. Each zone type follows a specific procedure detailed in the sections below:

- [Adding a Master Zone](#)
- [Adding a Slave Zone](#)
- [Adding a Forward Zone](#)
- [Adding a Stub Zone](#)
- [Adding a Hint Zone](#)
- [Adding a Delegation-Only Zone](#)
- SOLIDserver can be used for [Hosting Active Directory Domain Zones](#).

Note that you can also import zones, for more details refer to the section [Importing Zones](#).

If you want to create RPZ zones, refer to the chapter [DNS Firewall \(RPZ\)](#).

Adding a Master Zone

A master zone stocks the original zone important records for a certain name space and answers the other name servers queries regarding this space name.

The most common use of Master zones is to configure them with a slave zone of the same name that stores all its records, in a Master/Slave configuration where the master server contains the master zones and the slave server contains the slave zones. However, if you want to make the DNS available at all times, you can use a multiple master configuration. In which case, all DNS servers are master servers for each zone. This disposition requires to propagate any change, made to a zone file or the DNS configuration, to every DNS server configured as master. Therefore, we recommend that you manage your master servers through a DNS Multi-Master smart architecture. For more details regarding its configuration, refer to the section [Multi-Master Smart Architecture](#).

If you want to create an RPZ master name zone, refer to the chapter [DNS Firewall \(RPZ\)](#).

To add a master zone

1. In the sidebar, go to **MS DNS > Zones**. The page **All zones** opens.
2. In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
3. In the field **DNS server**, select the server of your choice.
4. Click on **[NEXT]**. The next page of the wizard opens.
5. If the server selected contains several views, the view selection page opens:
 - a. In the drop-down list **View**, select a view.
 - b. Click on **[NEXT]**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically and specified in the zone details on the next page of the wizard.

If your server does not contain any view, the page is not displayed.

6. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the list **DNS zone type**, select *Master*.
8. In the list **DNS zone resolution**: select *Name* or *Reverse* and click on **NEXT**. The next page of the wizard opens.
9. Configure the zone's basic parameters.

- a. For a **Name** zone:

Table 39.2. DNS name zone basic parameters

Field	Description
Name	Type in the zone name. It should strictly conform to the syntax given in RFC1034 ^a (page 7). If you are adding an Amazon Route 53 zone, you must include the zone TLD. This field is mandatory.
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc1034>.

- b. For a **Reverse** zone:

Table 39.3. DNS reverse zone basic parameters

Field	Description
Name	The name of the reverse domain, it auto-completes with the address you type in the next field. The suffix displayed changes according to the <i>Reverse type</i> selected.
IP address / IPv6 address	Type in the IP address for the zone. The address you type in completes the reverse domain name, it should be composed of a maximum of three bytes (xxx.xxx.xxx). This field is mandatory.
Reverse type	Select the reverse resolution method: <i>IPv4 in-addr.arpa</i> , <i>E164 arpa</i> , <i>IPv6 int</i> or <i>IPv6 arpa</i> . Once selected, the extension is automatically displayed in the field <i>Name</i> . This field is mandatory.
IPv4 in-addr.arpa	You can select this field to configure IPv4 reverse-mapping.
E164 arpa	You can select this field to configure telephone number mapping for the zone, it uses the phone numbers dedicated reverse mapping domain suffix (e164.arpa).
IPv6 int	You can select this field to configure IPv6 reverse-mapping. Note that this extension is deprecated, so unless your IPv6 configuration is older than 2001 we recommend that you use the IPv6 arpa extension. For more details, refer to RFC 4159 ^a .
IPv6 arpa	You can select this field to configure IPv6 reverse-mapping.
View	Select the view in which the zone should be created. If there are no views in the selected server, the list is empty.

Field	Description
Space	Select one of the IPAM spaces that should be tied to that zone or <i>None</i> . The selected space is updated by the DNS zone you are creating.

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc4159>.

- If you are managing an Active Directory integrated Microsoft DNS server, you might want to configure the **Expert Mode** and/or **AD integrated** parameters. If not, go to step 11.

Table 39.4. DNS zone expert mode parameters

Field	Description
Expert Mode	If your Microsoft DNS server is not managed via a smart, this box is not displayed. See next row. If you manage a Microsoft DNS server via a smart architecture, tick this box to complete the AD configuration. The box <i>AD integrated</i> appears.
AD integrated	If your server is AD integrated, tick this box to set your replication preferences. If you manage a Microsoft DNS server via a smart architecture, go to step 11. If your Microsoft DNS server is not managed via a smart, the drop-down list <i>AD Replication</i> appears.
AD replication	You can configure the zone content and parameters replication via the option: <i>All DC in the AD Domain (default)</i> , <i>All DNS servers in the AD domain</i> or <i>All DNS servers in the AD forest</i> . By default, the replication is set to <i>All DC in the AD Domain (default)</i> for zones created in a smart architecture. All DC in the AD domain: select this option to replicate the zone parameters and content to all the Domain Controllers of the AD domain. This option is selected by default, it is the only option available for zones created in a smart architecture. This option is not available for <i>Stub</i> zones. All DNS servers in the AD domain: select this option to replicate the zone parameters and content to all the DNS servers of the AD domain. All DNS servers in the AD forest: select this option to replicate the zone parameters and content to all the DNS servers of the AD forest.

Note that the column **AD integrated** on the page *All zones* allows to display the AD integration configuration (*Yes*, *No*, *N/A*) of each zone. For more details, refer to the section [Customizing the List Layout](#).

- Depending on the administrator configuration, you may be able to fill in the advanced properties fields following the table below:

Table 39.5. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The last page of the wizard opens.
- The fields on that page are compulsory and automatically filled. However you can edit them:

Table 39.6. DNS zone advanced parameters

Field	Description
Primary server	Specify the primary Master server for the zone. When you create a zone on a smart server, it is automatically filled and cannot be edited.
Responsible	Specify the administrator email address for the zone.
Serial number	The zone serial number. It is automatically incremented for each zone change.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expiration	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

You can set the value by default for the parameters above, except for the *Primary server* and *Serial number*. For more details, refer to the procedure [To configure the default SOA parameters of Master zones](#) below.

- Click on **OK** to complete the operation. The report opens and closes. The zone is listed and is marked *Delayed create* before being marked **OK**.

During the first Master zone addition, the allow-update option is by default configured with the ACL *admin*. Within SOLIDserver *admin* corresponds to *any*, so you might want to change the ACL and restrict the option use. For more details, refer to the section [Configuring DNS Update Authorizations on a Zone](#). Note that you can also configure such ACL for several zones at once. For more details, refer to the section [Setting Authorizations](#).

You can edit a master zone at any time. For more details, refer to the section [Editing a Master Zone](#).

Note that you can configure default values for the SOA record that is automatically created when you add a master zone.

To configure the default SOA parameters of Master zones

- Go to **DNS > Servers, Views or Zones**. The page opens.
- In the menu, select **Extra options > Default configuration**. The wizard **Default SOA parameters** opens.
- Edit the fields with the values you want to be automatically used when adding a master zone:

Table 39.7. DNS zone default SOA parameters

Field	Description
Responsible	Specify the administrator email address for the zone.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expiration	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.

Field	Description
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

- Click on **OK** to complete the operation. The report opens and closes. When adding a master zone, the DNS zone advanced parameters now auto-completes with the values you indicated for the SOA.

Adding a Slave Zone

A slave zone has one purpose, to respond to requests made to other servers that have authority over the domain queried.

They are usually created on slave name servers that receive their information from master name servers through a zone transfer, the master zone and the slave zone on each server are named the same. During the zone transfer, the master zone sends a NOTIFY to all the slave zone(s) it knows. The zone content is only sent to the slave zones authorized to receive the transfer, the other zones receive an error message. Note that several master servers can be configured for one slave server.

If you want to create an RPZ slave name zone, refer to the chapter [DNS Firewall \(RPZ\)](#).

To add a slave zone

- In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
- In the field **DNS server**, select the server of your choice.
- Click on **NEXT**. The next page of the wizard opens.
- If the server selected contains several views, the view selection page opens:
 - In the drop-down list **View**, select a view.
 - Click on **NEXT**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically.

If your server does not contain any view, the page is not displayed.

- If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the list **DNS zone type**, select *Slave*.
- In the list **DNS zone resolution**, select *Name* or *Reverse*.
- Click on **NEXT**. The next page of the wizard opens.
- Configure the zone's basic parameters.

- a. For a **Name** zone:

Table 39.8. DNS name zone basic parameters

Field	Description
Name	Type in the zone name. It should strictly conform to the syntax given in RFC1034 ^a (page 7). If you are adding an Amazon Route 53 zone, you must include the zone TLD. This field is mandatory.
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc1034>.

- b. For a **Reverse** zone:

Table 39.9. DNS reverse zone basic parameters

Field	Description
Name	The name of the reverse domain, it auto-completes with the address you type in the next field. The suffix displayed changes according to the <i>Reverse type</i> selected.
IP address / IPv6 address	Type in the IP address for the zone. The address you type in completes the reverse domain name, it should be composed of a maximum of three bytes (xxx.xxx.xxx). This field is mandatory.
Reverse type	Select the reverse resolution method: <i>IPv4 in-addr.arpa</i> , <i>E164 arpa</i> , <i>IPv6 int</i> or <i>IPv6 arpa</i> . Once selected, the extension is automatically displayed in the field <i>Name</i> . This field is mandatory.
IPv4 in-addr.arpa	You can select this field to configure IPv4 reverse-mapping.
E164 arpa	You can select this field to configure telephone number mapping for the zone, it uses the phone numbers dedicated reverse mapping domain suffix (e164.arpa).
IPv6 int	You can select this field to configure IPv6 reverse-mapping. Note that this extension is deprecated, so unless your IPv6 configuration is older than 2001 we recommend that you use the IPv6 arpa extension. For more details, refer to RFC 4159 ^a .
Ipv6 arpa	You can select this field to configure IPv6 reverse-mapping.
View	Select the view in which the zone should be created. If there are no views in the selected server, the list is empty.
Space	Select one of the IPAM spaces that should be tied to that zone or <i>None</i> . The selected space is updated by the DNS zone you are creating.

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc4159>.

11. Depending on the administrator configuration, you may be able to fill in the advanced properties fields:

Table 39.10. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .

Field	Description
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

12. Click on **NEXT**. The last page of the wizard opens.
13. Set up the list of master servers for the zone using the table below:

Table 39.11. DNS slave zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is mandatory. ^a
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

^aThe master zones of the server you specify through this IP address is automatically allowed to send notify messages of any changes to the slave zone you are creating.

Once the IP, port and key are configured, click on **ADD**. The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to **DELETE** or **UPDATE** it once created.

14. Click on **OK** to complete the operation. The report opens and closes. The zone is listed and is marked  *Delayed create* before being marked  **OK**.

Adding a Forward Zone

A forward zone, or *redirector*, allows to redirect all recursive requests for a zone toward a selected list of servers. The listed servers search local zones to resolve the recursive requests to which they cannot respond.

To add a forward zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
3. In the field **DNS server**, select the server of your choice.
4. Click on **NEXT**. The next page of the wizard opens.
5. If the server selected contains several views, the view selection page opens:
 - a. In the drop-down list **View**, select a view.
 - b. Click on **NEXT**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically.

If your server does not contain any view, the page is not displayed.

6. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the list **DNS zone type**, select *Forward*.
8. In the list **DNS zone resolution**, select *Name* or *Reverse*.
9. Click on **NEXT**. The next page of the wizard opens.
10. Configure the zone's basic parameters.

- a. For a **Name** zone:

Table 39.12. DNS name zone basic parameters

Field	Description
Name	Type in the zone name. It should strictly conform to the syntax given in RFC1034 ^a (page 7). If you are adding an Amazon Route 53 zone, you must include the zone TLD. This field is mandatory.
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc1034>.

- b. For a **Reverse** zone:

Table 39.13. DNS reverse zone basic parameters

Field	Description
Name	The name of the reverse domain, it auto-completes with the address you type in the next field. The suffix displayed changes according to the <i>Reverse type</i> selected.
IP address / IPv6 address	Type in the IP address for the zone. The address you type in completes the reverse domain name, it should be composed of a maximum of three bytes (xxx.xxx.xxx). This field is mandatory.
Reverse type	Select the reverse resolution method: <i>IPv4 in-addr.arpa</i> , <i>E164 arpa</i> , <i>IPv6 int</i> or <i>IPv6 arpa</i> . Once selected, the extension is automatically displayed in the field <i>Name</i> . This field is mandatory.
IPv4 in-addr.arpa	You can select this field to configure IPv4 reverse-mapping.
E164 arpa	You can select this field to configure telephone number mapping for the zone, it uses the phone numbers dedicated reverse mapping domain suffix (e164.arpa).
IPv6 int	You can select this field to configure IPv6 reverse-mapping. Note that this extension is deprecated, so unless your IPv6 configuration is older than 2001 we recommend that you use the IPv6 arpa extension. For more details, refer to RFC 4159 ^a .
IPv6 arpa	You can select this field to configure IPv6 reverse-mapping.
View	Select the view in which the zone should be created. If there are no views in the selected server, the list is empty.
Space	Select one of the IPAM spaces that should be tied to that zone or <i>None</i> . The selected space is updated by the DNS zone you are creating.

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc4159>.

11. If you are managing a Microsoft DNS server through a smart architecture, you might want to configure the parameters of the **Expert Mode**:

Table 39.14. DNS zone expert mode parameters

Field	Description
Expert Mode	If your Microsoft DNS server is not managed via a smart, this box is not displayed. See next row. If you manage a Microsoft DNS server via a smart architecture, tick this box to complete the AD configuration. The box <i>AD integrated</i> appears.
AD integrated	If your server is AD integrated, tick this box to set your replication preferences. If you manage a Microsoft DNS server via a smart architecture, go to step 11. If your Microsoft DNS server is not managed via a smart, the drop-down list <i>AD Replication</i> appears.
AD replication	You can configure the zone content and parameters replication via the option: <i>All DC in the AD Domain (default)</i> , <i>All DNS servers in the AD domain</i> or <i>All DNS servers in the AD forest</i> . By default, the replication is set to <i>All DC in the AD Domain (default)</i> for zones created in a smart architecture. All DC in the AD domain: select this option to replicate the zone parameters and content to all the Domain Controllers of the AD domain. This option is selected by default, it is the only option available for zones created in a smart architecture. This option is not available for <i>Stub</i> zones. All DNS servers in the AD domain: select this option to replicate the zone parameters and content to all the DNS servers of the AD domain. All DNS servers in the AD forest: select this option to replicate the zone parameters and content to all the DNS servers of the AD forest.

12. Depending on the administrator configuration, you may be able to fill in the advanced properties fields:

Table 39.15. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

13. Click on **NEXT**. The last page of the wizard opens.
14. Configure the forwarders and forward mode for the zone. The fields *Forwarders list* and *Forward Mode* are mandatory.

- a. In the field **Add a forwarder (IP)**, type in the IP address of the master server to which the queries should be forwarded and click on **ADD**.

The IP address is moved to the **Forwarders list**. Repeat these actions for as many servers as needed. The order of the servers in the list is not important.

- To update a forwarder, select an IP address in the list, its is displayed in in the field **Add a forwarder (IP)**: change the needed data and click on **UPDATE**.
- To delete a forwarder, select its IP address in the list and click on **DELETE**.

- b. Select a **Forward Mode** for the zone, it can be either *First*, *Only* or *None*. This field is mandatory.

Table 39.16. DNS forward zone parameters

Forward mode	Description
First	Select this option if you want the zone to first send a query to the forwarder, if not answered, it issues queries directly. This mode is selected by default.
Only ^a	Select this option if you only want the zone to forward queries.
None	Select this option if you configured a forward at server level in which case the forwarding is set by default.

^aIf you manage reverse forward zones for private or reserved addresses on a BIND server, you must select the forward mode *Only*.

Once the server IP address and the forward mode are configured, click on **+**. The configuration is listed in the **Forwarders list**.

- Click on **OK** to complete the operation. The report opens and closes. The zone is listed and is marked  *Delayed create* before being marked  **OK**.

Adding a Stub Zone

A stub zone is similar to a slave zone, with the exception that it does more than simply replicate the name servers of a master zone. Stub zones can be used to force the resolution of a domain, particularly for a restrained collection of servers. They are not part of the DNS standard zones, they are specific to BIND implementations.

The stub zone cannot be added into Generic DNS servers.

To add a stub zone

- In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
- In the field **DNS server**, select the server of your choice.
- Click on **NEXT**. The next page of the wizard opens.
- If the server selected contains several views, the view selection page opens:
 - In the drop-down list **View**, select a view.
 - Click on **NEXT**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically.

If your server does not contain any view, the page is not displayed.

- If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the list **DNS zone type**, select *Stub*.
- In the list **DNS zone resolution**, select *Name* or *Reverse*.
- Click on **NEXT**. The next page of the wizard opens.

10. Configure the zone's basic parameters.

- a. For a **Name** zone:

Table 39.17. DNS name zone basic parameters

Field	Description
Name	Type in the zone name. It should strictly conform to the syntax given in RFC1034 ^a (page 7). If you are adding an Amazon Route 53 zone, you must include the zone TLD. This field is mandatory.
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc1034>.

- b. For a **Reverse** zone:

Table 39.18. DNS reverse zone basic parameters

Field	Description
Name	The name of the reverse domain, it auto-completes with the address you type in the next field. The suffix displayed changes according to the <i>Reverse type</i> selected.
IP address / IPv6 address	Type in the IP address for the zone. The address you type in completes the reverse domain name, it should be composed of a maximum of three bytes (xxx.xxx.xxx). This field is mandatory.
Reverse type	Select the reverse resolution method: <i>IPv4 in-addr.arpa</i> , <i>E164 arpa</i> , <i>IPv6 int</i> or <i>IPv6 arpa</i> . Once selected, the extension is automatically displayed in the field <i>Name</i> . This field is mandatory.
IPv4 in-addr.arpa	You can select this field to configure IPv4 reverse-mapping.
E164 arpa	You can select this field to configure telephone number mapping for the zone, it uses the phone numbers dedicated reverse mapping domain suffix (e164.arpa).
IPv6 int	You can select this field to configure IPv6 reverse-mapping. Note that this extension is deprecated, so unless your IPv6 configuration is older than 2001 we recommend that you use the IPv6 arpa extension. For more details, refer to RFC 4159 ^a .
IPv6 arpa	You can select this field to configure IPv6 reverse-mapping.
View	Select the view in which the zone should be created. If there are no views in the selected server, the list is empty.
Space	Select one of the IPAM spaces that should be tied to that zone or <i>None</i> . The selected space is updated by the DNS zone you are creating.

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc4159>.

11. If you are managing a Microsoft DNS server through a smart architecture, you might want to configure the parameters of the **Expert Mode**:

Table 39.19. DNS zone expert mode parameters

Field	Description
Expert Mode	If your Microsoft DNS server is not managed via a smart, this box is not displayed. See next row. If you manage a Microsoft DNS server via a smart architecture, tick this box to complete the AD configuration. The box <i>AD integrated</i> appears.

Field	Description
AD integrated	If your server is AD integrated, tick this box to set your replication preferences. If you manage a Microsoft DNS server via a smart architecture, go to step 11. If your Microsoft DNS server is not managed via a smart, the drop-down list <i>AD Replication</i> appears.
AD replication	You can configure the zone content and parameters replication via the option: <i>All DC in the AD Domain (default)</i> , <i>All DNS servers in the AD domain</i> or <i>All DNS servers in the AD forest</i> . By default, the replication is set to <i>All DC in the AD Domain (default)</i> for zones created in a smart architecture. All DC in the AD domain: select this option to replicate the zone parameters and content to all the Domain Controllers of the AD domain. This option is selected by default, it is the only option available for zones created in a smart architecture. This option is not available for <i>Stub</i> zones. All DNS servers in the AD domain: select this option to replicate the zone parameters and content to all the DNS servers of the AD domain. All DNS servers in the AD forest: select this option to replicate the zone parameters and content to all the DNS servers of the AD forest.

- Depending on the administrator configuration, you may be able to fill in the advanced properties fields:

Table 39.20. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **NEXT**. The last page of the wizard opens.
- Set up the list of master servers for the zone using the table below:

Table 39.21. DNS stub zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is mandatory.
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

Once the IP, port and key are configured, click on **ADD**. The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to **DELETE** or **UPDATE** it once created.

- Click on **OK** to complete the operation. The report opens and closes. The zone is listed and is marked  *Delayed create* before being marked  **OK**. A stub zone only contains an SOA and NS RRs.

Adding a Hint Zone

A hint zone is aimed at querying root servers. By default, EfficientIP DNS servers embed a hint zone that is updated automatically but not listed on the page *All zones*. If you manage your DNS with another type of server, keep in mind that the hint zone is relevant only for name servers that provide recursive services.

The hint zone updates the local server cache with a list of the 13 root-servers saved in the form of A records (from *a.root-servers.net* to *m.root-servers.net*). Therefore, one hint zone per server or view is enough. When the server starts up, it uses the hint zone to query a root zone and obtain the complete list of the current authoritative root servers. This list is then used by the name server as a starting point for any domain query, if there is no locally defined zone (slave or master) or cached answers. A hint zone should be updated every 12 months or whenever there are discrepancies returned in the log messages, when the DNS server loads for instance.

Note that the hint zone can also contain an internal list and be used locally; in this case, the configuration is running an internal name service on a closed network, or the name server is not defined but recursive queries are required.

The hint zone cannot be added into a Microsoft server.

To add a hint zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
3. In the field **DNS server**, select the server on which you are adding a zone.
4. Click on **NEXT**. The next page of the wizard opens.
5. If the server selected contains several views, the view selection page opens:
 - a. In the drop-down list **View**, select a view.
 - b. Click on **NEXT**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically.

If your server does not contain any view, the page is not displayed.

6. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the list **DNS zone type**, select *Hint*.
8. In the list **DNS zone resolution**, select *Name* or *Reverse*.
9. Click on **NEXT**. The next page of the wizard opens.
10. Depending on the administrator configuration, you may be able to fill in the advanced properties fields:

Table 39.22. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on **OK** to complete the operation. The report opens and closes. The zone is listed, named and marked **OK**.

Adding a Delegation-Only Zone

The delegation-only zone allows to send an NXDOMAIN response to any query received without an explicit or implicit delegation in the authority section. They are usually created on caching/re-cursive servers.

When a zone is declared as delegation-only, it does not contain any record. You can use delegation-only zones to filter out wildcard or synthesized data from NAT boxes or from authoritative name servers whose undelegated (in-zone) data is of no interest. They can also be used to enforce the delegation-only status of infrastructure zones (e.g. COM, NET, ORG).

The name of the delegation-only zone is the domain for which you send the NXDOMAIN response, any subdomain responds normally.

The Delegation-Only zone cannot be added into a Microsoft server.

To add a delegation-only zone

- In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
- In the field **DNS server**, select the server on which you are adding a zone.
- Click on **NEXT**. The next page of the wizard opens.
- If the server selected contains several views, the view selection page opens:
 - In the drop-down list **View**, select a view.
 - Click on **NEXT**. The next page of the wizard opens.

If your server contains only one view, the page is not displayed. Your view is selected automatically.

If your server does not contain any view, the page is not displayed.

- If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the list **DNS zone type**, select *Delegation-Only*.
8. In the list **DNS zone resolution**, select *Name* or *Reverse*.
9. Click on **NEXT**. The next page of the wizard opens.
10. Configure the zone's basic parameters.
 - a. For a **Name** zone:

Table 39.23. DNS name zone basic parameters

Field	Description
Name	Type in the zone name. It should strictly conform to the syntax given in RFC1034 ^a (page 7). If you are adding an Amazon Route 53 zone, you must include the zone TLD. This field is mandatory.
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc1034>.

- b. For a **Reverse** zone:

Table 39.24. DNS reverse zone basic parameters

Field	Description
Name	The name of the reverse domain, it auto-completes with the address you type in the next field. The suffix displayed changes according to the <i>Reverse type</i> selected.
IP address / IPv6 address	Type in the IP address for the zone. The address you type in completes the reverse domain name, it should be composed of a maximum of three bytes (xxx.xxx.xxx). This field is mandatory.
Reverse type	Select the reverse resolution method: <i>IPv4 in-addr.arpa</i> , <i>E164 arpa</i> , <i>IPv6 int</i> or <i>IPv6 arpa</i> . Once selected, the extension is automatically displayed in the field <i>Name</i> . This field is mandatory.
IPv4 in-addr.arpa	You can select this field to configure IPv4 reverse-mapping.
E164 arpa	You can select this field to configure telephone number mapping for the zone, it uses the phone numbers dedicated reverse mapping domain suffix (e164.arpa).
IPv6 int	You can select this field to configure IPv6 reverse-mapping. Note that this extension is deprecated, so unless your IPv6 configuration is older than 2001 we recommend that you use the IPv6 arpa extension. For more details, refer to RFC 4159 ^a .
Ipv6 arpa	You can select this field to configure IPv6 reverse-mapping.
View	Select the view in which the zone should be created. If there are no views in the selected server, the list is empty.
Space	Select one of the IPAM spaces that should be tied to that zone or <i>None</i> . The selected space is updated by the DNS zone you are creating.

^aAvailable on IETF website at <http://tools.ietf.org/html/rfc4159>.

Keep in mind that the name of your delegation-only zone is a domain that, once queried, sends an NXDOMAIN response to the client.

11. Depending on the administrator configuration, you may be able to fill in the advanced properties fields:

Table 39.25. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes. The zone is listed and marked  **OK**.

Hosting Active Directory Domain Zones

SOLIDserver can be configured to update native Microsoft DNS servers and host DNS zones coming from AD domains.

On Active Directory integrated Microsoft DNS servers, when you add *Master*, *Slave*, *Forward* or *Stub* zones the box **AD integrated** allows to set up the AD replication of your choice. Note that once set, you cannot edit the AD replication configuration unless you delete the zone and recreate it. For more details, refer to the sections [Adding a Master Zone](#), [Adding a Slave Zone](#), [Adding a Forward Zone](#) or [Adding a Stub Zone](#).

Keep in mind that the **DNS Multi-Master smart architecture can reproduce Microsoft's Multi-Master behavior**. For more details, refer to the section [Multi-Master Smart Architecture](#).

Synchronizing DNS Zones

It is possible to refresh the management database with the content of one or more zones that have not been edited from SOLIDserver. This synchronization is done automatically after a period defined by the refresh parameter of the zone (SOA), but the administrator can force a synchronization at any time to update the zone data.

To retrieve the entire content of a zone, you can use the option *Force full synchronization*. For more details, refer to the section [Forcing SOLIDserver to Retrieve the Full Content of a Zone](#).

To synchronize zones

- In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
- Tick the zone(s) you want to synchronize.
- In the menu, select  **Edit > Status > Synchronize**. The wizard **Synchronization** opens.
- Click on to complete the operation. The report opens and closes when the synchronization is over. The page reloads.

Forcing SOLIDserver to Retrieve the Full Content of a Zone

The option *Force full synchronization* resets the serial number of the SOA of a zone. This action allows SOLIDserver to retrieve and upload the complete content of a zone managed from the

GUI, contrary to a regular synchronization that performs an incremental retrieval of a zone records which only retrieve the latest changes.

Resetting the SOA serial ensures that the entire zone data is up to date. For instance, this option is useful if you restored a server that you used to manage, as you might have a synchronization drift between the data that SOLIDserver stored locally and all the changes that have been performed directly on the server since you last managed it.

To force the full synchronization for your zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) you want to synchronize.
3. In the menu, select **Tools > Expert > Force full synchronization**. The wizard **Synchronization** opens.
4. Click on **OK** to complete the operation. The report opens and closes when the synchronization is over. The page reloads.

Editing DNS Zones

SOLIDserver allows you to edit all zone types.

If you want to edit RPZ zones, refer to the chapter [DNS Firewall \(RPZ\)](#).

Editing a Master Zone

After [Adding a Master Zone](#), you can always edit the zone configuration parameters from its properties page or from the page *All zones* itself, using **✎** in the contextual menu.

You cannot edit the *AD replication* configuration of the zones of an Active Directory integrated Microsoft DNS server.

To edit a master zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **✎**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS zone** opens.

The top gray area sums up the zone parameters: its *Name*, *Type*, *Server* and *View*.

4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Space** and **Advanced properties** if need be:

Table 39.26. Space and Advanced properties fields

Field	Description
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated

Field	Description
	advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

Note that you cannot edit any other parameter, including the configuration of the parameters *AD integrated* and *AD replication*.

- Click on **NEXT**. The last page of the wizard opens.
- Edit the advanced parameters if need be:

Table 39.27. DNS zone advanced parameters

Field	Description
Primary server	Specify the primary Master server for the zone. When you create a zone on a smart server, it is automatically filled and cannot be edited.
Responsible	Specify the administrator email address for the zone.
Serial number	The zone serial number. It is automatically incremented for each zone change.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expiration	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

- Click on **OK** to complete the operation. The report opens and closes. The changes are visible in the panel **Main properties**.

You can also add and delete NS records, for multiple master zones at once, from the page *All zones*.

Editing a Slave Zone

After [Adding a Slave Zone](#), you can always edit the zone configuration parameters from its properties page or from the page *All zones* itself, using  in the contextual menu.

You cannot edit the *AD replication* configuration of the zones of an Active Directory integrated Microsoft DNS server.

To edit a slave zone

- In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.

2. At the end of the line of the zone of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on [\[EDIT\]](#). The wizard **Edit a DNS zone** opens.

The top gray area sums up the zone parameters: its *Name*, *Type*, *Server* and *View*.

4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on [\[NEXT\]](#). The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Space** and **Advanced properties** if need be:

Table 39.28. Space and Advanced properties fields

Field	Description
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

Note that you cannot edit any other parameter, including the configuration of the parameters *AD integrated* and *AD replication*.

6. Click on [\[NEXT\]](#). The last page of the wizard opens.
7. If you want to add another master server, refer to the step 14 of the procedure [To add a slave zone](#).
8. If you want to edit a server, select it in the list **Masters**, the parameters configured appear in the fields **Master IP address**, **Port** and **TSIG key**: modify the content of any field according to your needs and click on [\[UPDATE\]](#). The server is modified in the list.
9. If you want to delete a server, select it in the list **Masters** and click on [\[DELETE\]](#). The server is no longer listed in the list.
10. Click on [\[OK\]](#) to complete the operation. The report opens and closes. The changes are visible in the panel **Main properties**.

Editing a Forward Zone

After [Adding a Forward Zone](#), you can always edit the zone from its properties page or from the page *All zones* itself, using  in the contextual menu.

Keep in mind that:

- The panel *Main properties* of a forward zone only allows to edit the zone classes and the advanced properties, if relevant.

- The panel *Forwarding* allows to edit the forwarding configuration, as detailed in the procedure below.
- You cannot edit the *AD replication* configuration of the zones of an Active Directory integrated Microsoft DNS server.

To edit a forward zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Forwarding**, click on **[EDIT]**. The wizard **Edit a DNS zone** opens.
4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Space** and **Advanced properties** if need be:

Table 39.29. Space and Advanced properties fields

Field	Description
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

Note that you cannot edit any other parameter, including the configuration of the parameters *AD integrated* and *AD replication*.

6. Click on **[NEXT]**. The last page of the wizard opens.
7. If you want to add another forwarding master server refer to the step 14 of the [To add a forward zone](#) procedure.
8. In the fields **Add a forwarder (IP)** and **Forward Mode**, fill in the address of the master server and select if the zone should forward *Only* or send a query *First*.
9. If you want to delete a server, select it in the list **Forwarders list** and click on **⌵**. The server is no longer listed in the list.
10. Click on **[OK]** to complete the operation. The report opens and closes. The changes are visible in the panels **Main properties** and **Forwarding**.

Editing a Stub Zone

After [Adding a Stub Zone](#), you can always edit the zone configuration parameters from its properties page or from the page *All zones* itself, using **⌵** in the contextual menu.

To edit a stub zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DNS zone** opens.
4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Space** and **Advanced properties** if need be:

Table 39.30. Space and Advanced properties fields

Field	Description
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

Note that you cannot edit any other parameter, including the configuration of the parameters *AD integrated* and *AD replication*.

6. Click on **[NEXT]**. The last page of the wizard opens.
7. If you want to add another master server refer to the step 14 of the [To add a stub zone](#) procedure.
8. If you want to edit a server, select it in the list **Masters**, the parameters configured appear in the fields **Master IP address**, **Port** and **TSIG key**: modify the content of any field according to your needs and click on **[UPDATE]**. The server is modified in the list.
9. If you want to delete a server, select it in the list **Masters** and click on **[DELETE]**. The server is no longer listed in the list.
10. Click on **[OK]** to complete the operation. The report opens and closes. The changes are visible in the panel **Main properties**.

Editing a Hint Zone

After [Adding a Hint Zone](#), you can always edit the zone from its properties page or from the page *All zones* itself, using **⌵** in the contextual menu.

Keep in mind that you can only edit the advanced properties of a hint zone.

To edit a hint zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DNS zone** opens.
4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Advanced properties** if need be:

Table 39.31. Advanced properties configuration

Field	Description
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

6. Click on **[OK]** to complete the operation. The report opens and closes. The changes are visible in the panel **Advanced properties**.

Editing a Delegation-Only Zone

After [Adding a Delegation-Only Zone](#), you can always edit the zone configuration parameters from its properties page or from the page *All zones* itself, using **⌵** in the contextual menu.

You cannot edit the *AD replication* configuration of the zones of an Active Directory integrated Microsoft DNS server.

To edit the properties of a delegation-only zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a DNS zone** opens.
4. If you or your administrator created classes at the zone level, in the list **DNS zone class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the **Space** and **Advanced properties** if need be:

Table 39.32. Space and Advanced properties fields

Field	Description
Space	You can select an existing IPAM space to associate with the zone or <i>None</i> . The selected space is updated by the zone records if you configure the dedicated advanced properties. For more details, refer to the section Configuring DNS Advanced Properties .
Advanced properties	Select either <i>Default</i> or <i>All</i> .
Default	This value indicates that all the fields/options displayed in the wizard are part of the wizard default display. This display configuration is available to your administrator via the wizard <i>Advanced properties customization</i> . For more details, refer to the chapter Managing Advanced Properties ; the DNS properties are detailed in the section Configuring DNS Advanced Properties .
All	This value allows to display all the advanced property fields/options that can be displayed in this wizard. For more details, refer to the chapter Managing Advanced Properties .

- Click on to complete the operation. The report opens and closes, the properties page refreshes. The changes are visible in the panels **Main properties** and **Advanced properties**.

Converting DNS Zones

When managing EfficientIP DNS servers, you can convert multiple DNS zones at once from master to slave and vice versa. The other zone types cannot be converted.

Note that:

- During a conversion from slave to master, you can choose to delete the old NS records it contains.
- Converting master zones to slave on servers that do not support slave zones (e.g. Route53 and Unbound) deletes all the records of the original master zone.

To convert slave zones to master

- In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- Tick the slave zone(s) you want to convert.
- In the menu, select **Edit > Convert > To master**. The wizard **Zone conversion from slave to master** opens.
- In the drop-down list **Remove old NS** you can decide what to do with the zone's NS records:
 - Select *Yes* to delete the NS record of the former slave zone.
 - Select *No* to keep them in the master zone for future use.
- Click on to complete the operation. The report opens and closes. The page reloads. The converted zone is displayed on the page.

To convert master zones to slave

- In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- Tick the master zone(s) you want to convert.

3. In the menu, select  **Edit > Convert > To slave**. The wizard **Zone conversion from master to slave** opens.
4. In the field **IP of master server**, specify the IP address of the master server that now has authority over the zone(s).
5. Click on  to complete the operation. The report opens and closes. The page reloads. The converted zone is displayed on the page.

Adding or Removing an NS Record

You can add and delete the same NS RR, at once, for multiple Master and Hint zones.

To add or remove an NS RR for multiple zones

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. Tick the Master or Hint zone(s) for which you want to add/delete an NS RR.
3. In the menu, select  **Edit > Add/Delete an NS RR**. The wizard **Add/Delete an NS RR** opens.
4. If you want to remove an NS RR from the selected zone(s), in the field **NS value to delete** type in the value of the record.
5. If you want add an NS RR to the selected zone(s), in the field **NS value to add** type in the value of the record.
6. Click on  to complete the operation. The report opens and closes. The page **All zones** reloads.
7. In the column **Name**, click on the name of the zone(s) you edited to display the records it now contains.

Copying or Moving DNS Zones

At some point you might need to migrate or copy zones from one DNS server or view to the other. In this case, you need to use the *Migrate* option. Note that this option has nothing to do with the zones database replication of the DNS command *allow-transfer*.

Copying or moving a zone also applies to the records it manages.

To copy a zone

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. Tick the zone (s) you want to copy on another server or view.
3. In the menu, select  **Edit > Migrate**. The wizard **Copying/Moving zones** opens.
4. In the drop-down list **Method**, select *Copy*.
5. In the drop-down list **Target server**, select the DNS server where you want to copy the selected zone. The wizard refreshes.
6. If the selected server has views, the drop-down list **Target view** appears, select the view of your choice. The wizard refreshes.
7. Tick the box **Asynchronous** to run the creation of records in the background. This option shortens the process but the records do not appear instantly on the page *All RRs*.

8. Click on to complete the operation. The report opens and closes. The page **All zones** is visible again and displays the duplicated zone. If you selected a view, the zone is also listed in the list **All zones** of said view.

To move a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone (s) you want to copy on another server or view.
3. In the menu, select **Edit > Migrate**. The wizard **Copying/Moving zones** opens.
4. In the drop-down list **Method**, select *Move*.
5. In the drop-down list **Target server**, select the DNS server where you want to move the selected zone. The wizard refreshes.
6. If the selected server has views, the **Target view** drop-down list appears, select the view of your choice. The wizard refreshes.
7. Tick the box **Asynchronous** to run the creation of records in the background. This option shortens the process but the records do not appear instantly on the page *All RRs*.
8. Click on to complete the operation. The report opens and closes. The page **All zones** is visible again and displays the migrated zone. If you selected a view, the zone is also listed on the page **All zones** of said view.

Setting Properties on Multiple DNS Zones

You can configure properties for multiple zones at once from the page *All zones*. This allows to set common properties for several zones, among these properties you can set IPAM spaces, forwarders, master servers or configure authorizations for specific users. Keep in mind that you can define new settings or use the parameters of existing zones.

Setting a Space

You can associate multiple zones, of any type, with the same IPAM space.

To set a space for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set a space.
3. In the menu, select **Edit > Properties > Set space**. The wizard **Set a Space** opens.
4. In the drop-down list **Space**, select the space of your choice.
5. Click on to complete the operation. The report opens and closes. On the properties page of the zone(s), the space is visible in the panel **Main properties**.

Setting Authorizations

You can set different types of authorizations and restrictions at zone level to [limit zone transfers](#) and [queries](#) or [manage DNS update](#). You can set the same authorizations properties for multiple zones at once.

Keep in mind that:

- You can configure the *allow-transfer* authorization only for master and slave zones.

- You can configure the *allow-query* authorization only for master, slave and stub zones.
- You can configure the *allow-update* authorization only for master zones.

To set allow-transfer properties for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set allow-transfer or allow-query properties.
3. In the menu, select **Edit > Properties > Set authorizations > Allow-transfer**. The wizard **Set zone transfer authorizations** opens.
4. To set new authorization parameters for the selected zone(s):
 - a. In the drop-down list **Source**, select *New settings*.
 - b. Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 39.33. Restriction and permission parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

5. To use existing authorization parameters for the selected zone(s):
 - a. In the drop-down list **Source**, select *Use existing zone parameters*. The wizard refreshes and the drop-down list *Zone* appears.
 - b. In the drop-down list **Zone**, select the zone whose properties you want to apply to your selection.
6. Click on **OK** to complete the operation. The report opens and closes. The changes are visible on the zone(s) properties page, in the panel **Main properties**.

To set allow-query properties for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set allow-transfer or allow-query properties.
3. In the menu, select **Edit > Properties > Set authorizations > Allow-query**. The wizard **Set zone queries authorizations** opens.
4. To set new authorization parameters for the selected zone(s):
 - a. In the drop-down list **Source**, select *New settings*.
 - b. Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 39.34. Restriction and permission parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <code><ip-address>/<prefix></code> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost and localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

5. To use existing authorization parameters for the selected zone(s):
 - a. In the drop-down list **Source**, select *Use existing zone parameters*. The wizard refreshes and the drop-down list *Zone* appears.
 - b. In the drop-down list **Zone**, select the zone whose properties you want to apply to your selection.
6. Click on **OK** to complete the operation. The report opens and closes. The changes are visible on the zone(s) properties page, in the panel **Main properties**.

To set allow-update properties for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set allow-update properties.

3. In the menu, select **Edit > Properties > Set authorizations > Allow-update**. The wizard **Set zone update authorizations** opens.
4. In the field **Source**, you can choose:

To configure a new update policy for your zone, following step 6.

To configure an update policy using existing zone parameters, following step 7.

To configure an update policy using GSS-TSIG, following step 8.

5. If you want to configure a new update policy parameters for the selected zone(s):
 - a. In the drop-down list **Source**, select *New settings*.
 - b. Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 39.35. Restriction and permission parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <code><ip-address>/<prefix></code> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin^a, any, none, localhost and localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

6. If you want to set the properties of an existing zone to the selected zone(s):
 - a. In the drop-down list **Source**, select *Use existing zone parameters*. The wizard refreshes and the drop-down list *Zone* appears.
 - b. In the drop-down list **Zone**, select the zone whose properties you want to apply to your selection.
7. If you want to set a specific update policy:
 - a. In the drop-down list **Source**, select *New settings*.
 - b. Tick the box *Use GSS-TSIG/update-policy*. The wizard refreshes, the update-policy related fields replace the other fields of the wizard.

- c. Configure the dynamic update permissions and restrictions for your zone:

Table 39.36. Update-policy configuration parameters

Option	Description
Permission	Specify if you <i>grant</i> or <i>deny</i> update-policy rights. This field applies to all the other fields detailed below.
Identity	Specify the GSS-TSIG key sent by the client when they try to update the zone. It may look like <code><ad-server-name>\$_@<full-ad-domain></code> . Its use depends on the <i>Matchtype</i> you select, it can therefore partially depend on the <i>Tname</i> .
Matchtype	Select the matchtype that suits your needs: <i>subdomain</i> , <i>krb5-self</i> , <i>krb5-subdomain</i> , <i>ms-self</i> or <i>ms-subdomain</i> , described below. The matchtype only applies to the record type specified in the list Any .
	subdomain allows to define the subdomain being updated. You must specify the domain in the field Tname following the format <code><domain>.<tld></code> , everything left of the specified <code><domain></code> becomes a match.
	wildcard allows to define the record being updated. You must specify the record's name in the field Tname using at least one wildcard (*). It can contain only a wildcard, in which case any record name can match.
	krb5-self allows to define a rule based on Kerberos machine principal (host/QDN@REALM). The record being updated matches the QDN part of the Principal. The matching REALM must be specified exactly in the fields Identity and Tname .
	krb5-subdomain allows to define a rule based on Kerberos machine principal (host/QDN@REALM). The subdomain being updated matches the QDN part of the Principal. The matching REALM is what is specified the field Identity , or any subdomain of the specified Identity .
	ms-self allows to define a rule based on AD format principal (machinename\$@REALM) to update machinename.realm in the DNS. The matching REALM must be specified exactly in the fields Identity and Tname .
	ms-subdomain allows to define a rule based on AD format principal (machine-name\$@REALM) to update machinename.realm in the DNS. The matching REALM is what is specified the field Identity , or any subdomain of the specified Identity .
Tname	Specify a value to which applies the Matchtype <i>subdomain</i> , <i>krb5-self</i> , <i>krb5-subdomain</i> , <i>ms-self</i> or <i>ms-subdomain</i> . The expected format is detailed above, with each Matchtype.
RR type	Select which record type(s) the configuration set in the previous fields applies to. It can be <i>Specific</i> or <i>Any</i> .
<i>Specific</i>	Configure permissions for the record types of your choice via the lists Available types and Selected type(s) . In the list Available types , select <i>A</i> , <i>AAAA</i> , <i>CNAME</i> , <i>HINFO</i> , <i>AFSDB</i> , <i>MX</i> , <i>PTR</i> , <i>NS</i> , <i>SRV</i> , <i>TXT</i> , <i>WKS</i> , <i>NSAP</i> or <i>DNAME</i> and click on <input type="checkbox"/> to move it to the list Selected type(s) . Repeat this action for as many record types as you need.
<i>Any</i>	Allows to apply the configuration to all the update-policy record types: <i>A</i> , <i>AAAA</i> , <i>CNAME</i> , <i>HINFO</i> , <i>AFSDB</i> , <i>MX</i> , <i>PTR</i> , <i>NS</i> , <i>SRV</i> , <i>TXT</i> , <i>WKS</i> , <i>NSAP</i> and <i>DNAME</i> .

- d. Once you configured the fields, click on . Your update-policy entry is moved to the **Update-policy list**. The page refreshes.
- e. You can configure as many entries as you want.

To organize the list, use the buttons and . Each restriction or permission is reviewed and processed following the order set in the list.

8. Click on to complete the operation. The report opens and closes. The changes are visible on the zone(s) properties page, in the panel **Main properties**.

Setting Forwarders

You can associate multiple master, slave, forward and stub zones to the same forward parameters at once.

To set forwarders for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set properties.
3. In the menu, select **Edit > Properties > Set forwarders**. The related wizard opens.
4. If you want to set the properties of an existing zone to the selected zone(s):
 - a. In the drop-down list **Source**, select *Use existing zone parameters*. The wizard refreshes and the drop-down list *Zone* appears.
 - b. In the drop-down list **Zone**, select the zone whose properties you want to apply to your selection.
5. If you want to set new properties to the selected zone(s):
 - a. In the drop-down list **Source**, select *New settings*.
 - b. In the drop-down list **Forward mode**, select *Default*, *None*, *First* or *Only* according to your needs. If you select *Only* or *First*, the field *Add a forwarder (IP)* appears.
 - c. If you selected *First* or *Only*, set your forwarders list:
 - i. In the field **Add a forwarder (IP)**, type in the address of a forwarder.
 - ii. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.
6. Click on **ADD**. The configuration is displayed in the list *ACL* at the bottom.
7. Repeat the operation with as many configuration parameters as needed.
8. Click on **OK** to complete the operation. The report opens and closes. The changes are visible on the zone(s) properties page, in the panel **Main properties**.

Setting Master Servers

You can associate multiple slave and stub zones to the same master server(s) at once.

To set a master server for multiple zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) for which you want to set properties.
3. In the menu, select **Edit > Properties > Set master servers**. The wizard **Set several master servers** opens.
4. In the field **Master IP address**, type in the IP address of the master server that communicates with the selected slave zone(s).
5. In the field **Port**, specify a port number on the master server dedicated to the communication with the selected slave zone(s).

6. In the field **TSIG key**, type in the TSIG key chosen to identify the zone(s) from the master server.
7. Click on **ADD**. The configuration is displayed in the list *ACL* at the bottom.
8. Repeat the operation with as many configuration parameters as master servers.
9. Click on **OK** to complete the operation. The report opens and closes. The changes are visible on the zone(s) properties page, in the panel **Main properties**.

Forcing the Update of DNS Zones' Content

By default, the zone content is updated every hour. If you know changes were made locally on the zones of the master or slave servers you manage, you can update the content of one or several zones at a time.

You can force a notify, refresh or retransfer right away if needed. Keep in mind that:

- You can only force a zone content update on BIND and EfficientIP DNS servers. Other servers do not support RNDG commands
- You cannot use these options on Hybrid DNS servers. For more details, refer to the chapter [Hybrid DNS Service](#).

Forcing DNS Zone Notification

When the records of a zone are edited, slave servers have to be notified of this change to ask for a zone transfer. Notify messages can be sent to name servers and/or IP address(es). SOLIDserver automatically triggers the notification when you are [Configuring DNS Notify Messages at Zone Level](#), but you can force the notification for master and slave zones at any time.

The option Force notify allows to send a notify, or also-notify if it is configured, on the zone(s) you select. It uses the information of NS records named like the zone you selected.

Keep in mind that the option cannot work if the zone selected does not contain at least one NS record named after the zone itself, or if the Notify is disabled on the zone,

To force master zones notification

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s), slave or master, that you want to send a notify from. The notify is sent only if the zone contains NS records named like the zone.
3. In the menu, select **Edit > Command > Force notify**. The wizard **Force zone notification** opens.
4. Click on **OK** to complete the operation. The report opens and closes when the operation is over. The page reloads.

Forcing DNS Zones Refresh

When the records of a master zone are edited, you can force the related slave zone(s) to retrieve the new values. Records from the master zone(s) that are not present in the slave zone(s) you want to refresh are not added. To force their creation, refer to the section [Forcing DNS Zones Retransfer](#).

To force slave zones refresh

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the slave zone(s) for which you want the master zone to refresh the records.
3. In the menu, select **Edit > Command > Force refresh**. The wizard **Force zone refresh** opens.
4. Click on **OK** to complete the operation. The report opens and closes when the operation is over. The page reloads.

Forcing DNS Zones Retransfer

When the records of a master zone are edited and new records are added, you can force the related slave zone(s) to retrieve all the new values and records it now contains. This operation triggers the creation and update, within the selected slave zone(s), of both the records present in the master zone(s) and their values.

If your slave zone contained more records than the master zone, the option *force retransfer* deletes the outdated records. Keep in mind that these changes are performed right away in the DNS database, but you need to synchronize the zone(s) after a *force retransfer* to see the changes in the GUI. For more details, refer to the section [Synchronizing DNS Zones](#).

To retrieve only the values of the records contained in the slave zone(s), but not the new records, refer to the section [Forcing DNS Zones Refresh](#).

To force slave zones retransfer

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Force the retransfer of the zone(s) your choice:
 - a. Tick the slave zone(s) for which you want the master zone to retransfer the records.
 - b. In the menu, select **Edit > Command > Force retransfer**. The wizard **Force zone retransfer** opens.
 - c. Click on **OK** to complete the operation. The report opens and closes when the operation is over. The page reloads.
3. Synchronize the zone(s) to see the changes in the GUI:
 - a. Tick the zone(s) you want to synchronize.
 - b. In the menu, select **Edit > Status > Synchronize**. The wizard **Synchronization** opens.
 - c. Click on **OK** to complete the operation. The report opens and closes when the synchronization is over. The page reloads.

Disabling and Enabling DNS Zones

All existing zones can be enabled and disabled from the management console, providing a viable option for stopping the availability of zones on one or several servers in one operation. This option is especially helpful when you have to move or repair servers for particular zones. When you disable a zone, its status switches to **Unmanaged**.

To enable/disable a zone

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. Tick the zones you want to disable.
3. In the menu, select  **Edit > Status > Enable** or **Disable**. The wizard opens.
4. Click on  to complete the operation. The report opens and closes. The zone status changes to  **OK** or  **Unmanaged**.

Deleting DNS Zones

The deletion procedure is the same for every type of zones. Deleting a zone also deletes all the resource records it contains.

To delete a zone

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. Filter the list if need be.
3. Tick the zone(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on  to complete the operation. The report opens and closes. The zone is marked  **Delayed delete** until it is no longer listed.

Defining a DNS Zone as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a zone as one of the resources of a specific group allows the users of that group to manage the zone in question as long as are they are granted the relevant rights.

Granting access to a zone as a resource also makes every item it contains available. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 40. Configuring DNS Zones

Like servers and views, zones can be configured individually to set a series of behaviors for the records they contain. Any configuration set at zone level overwrites what was set at server (whether physical or smart) and view level.

Managing DNS Zone Delegation

DNS provides the option of dividing up the namespace into one or more zones, which can then be stored, distributed and replicated to other DNS servers. When considering dividing your DNS namespace to make additional zones, take into account the following reasons to use additional zones:

- A need to delegate management of part of your DNS namespace to another location or department within your organization.
- A need to divide one large zone into smaller zones to distribute traffic loads among multiple servers, improve DNS name resolution performance, or create a more fault-tolerant DNS environment.
- A need to extend the namespace by adding numerous subdomains at once, to accommodate the opening of a new branch or site for instance.

If, for any of these reasons, you could benefit from delegating zones, it might make sense to restructure your namespace by adding additional zones. When choosing how to structure zones, you should use a plan that reflects the structure of your organization. When delegating zones within your namespace, be aware that for each new zone you create, you need delegation records (NS) in other zones that point to the authoritative DNS servers for the new zone. This is necessary both to transfer authority and to provide correct referral to other DNS servers and clients of the new servers being made authoritative for the new zone.

To make a server known to others outside of the new delegated zone, two RRs are needed in the parent zone to complete delegation to the new zone. These RRs include:

- An NS RR to effect the delegation. This RR is used to advertise that the server named is an authoritative server for the delegated subdomain.
- An A RR (also known as a glue record) is needed to resolve the name of the server specified in the NS RR to its IP address. The process of resolving the host name in this RR to the delegated DNS server in the NS RR is sometimes referred to as glue chasing. In reality, the A record is not required when it comes to configuring zones delegation; however, if you add it, you save the DNS client some time as you give in one query the authoritative server of the child zone and IP address. That way, there is no need to query twice to first get the server and then its IP address.

Configuring Delegation at Zone Level

At zone level, setting up the delegation implies editing the properties of the zone with the appropriate data.

To configure a name server for a zone

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on . The properties page opens.

3. In the panel **Name Servers**, click on **[EDIT]**. The wizard **Authoritative DNS servers** opens.
4. In the field **DNS server**, type in the name of the server of your choice. Repeat these actions for as many servers as needed.
5. Click on **[ADD]**. The server is moved to the list **Authoritative DNS servers**.
6. Click on **[NEXT]**. The page **Delegated data** opens.
7. In the field **Delegation**, type in the name of the RR and the server you want to delegate it following the syntax: *RRname > dnsserver.name*.
8. Click on **[ADD]**. Your data is moved to the list **Delegated data**. Repeat these actions for as many RRs and servers as needed.
9. Click on **[OK]** to complete the operation. The report opens and closes. The configuration parameters are visible in the panel.

Configuring delegation only creates the NS record. For more details regarding the A record addition, refer to the section [Configuring the Delegation at Record Level](#).

Using the Classless in-addr.arpa Delegation

SOLIDserver allows you to configure a classless in-addr.arpa delegation for networks containing less than 256 IP addresses, as defined in the RFC 2317. Typically, it is used to delegate reverse DNS lookup for part of that network to other DNS servers.

In the parent master reverse zone, the classless in-addr.arpa delegation creates CNAME resource records for each address you want to delegate. It also creates an NS RR for each delegated server. Note that the NS record of each delegated server can be created in a domain different from *in-addr.arpa* using a suffix for the CNAME RRs value. For the reverse lookup to function properly, the delegated server(s) should contain the PTR records associated to each address.

To add a classless in-addr.arpa delegation

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Name**, click on the master reverse zone you want to delegate a part of. It should be composed of a maximum of three bytes (*xxx.xxx.xxx.in-addr.arpa*). The page **All RRs** opens and displays the RRs of the zone.
3. In the menu, select **+ Add > Classless in-addr.arpa delegation**. The wizard **Add a classless in-addr.arpa delegation** opens.
4. In the field **Start address**, type in the first address of the range you want to delegate. By default, the first available address of the zone id displayed in this field.
5. In the field **Delegation range size**, type in the number of addresses you want to delegate.
6. In the field **Delegated NS**, type in the name of the DNS server) that should be authoritative over the range of addresses. Use **[+]** to add this server name the **Delegated NS list**. Repeat these actions for as many servers as needed. Use **[-]** to remove a server name from the list.
7. In the drop-down list **Delegated zone format**, select the concatenation format (*[start]-[end].c.b.a.in-addr.arpa*, *[start]-[size].c.b.a.in-addr.arpa*, *[start]-[prefix].c.b.a.in-addr.arpa*) for NS RR name.
8. Tick the box **Add a specific suffix** if you want the NS RR to be created in a domain different from *in-addr.arpa*. The **Specific suffix** field appears.

9. In the field **Specific suffix**, type in the suffix of your choice. This suffix corresponds to the domain in which you want to create the NS RR. This suffix is added at the end of each of the CNAME RR you are creating.
10. Click on **OK** to complete the operation. The report opens and closes. The page **All RRs** is visible again. There are as many CNAME RRs as delegated addresses and as many a NS records as delegated servers. In the column **Value**, each address is listed according to the format you chose, if you added a suffix, it is visible in that column as well.

Configuring DNS Forwarding at Zone Level

You can set a forwarding configuration on master, slave, forward and stub zones from their properties page. Note that:

- On a zone belonging to a server not managed via a smart architecture, the specific forwarding configuration only applies to the zone on said server.
- On a zone belonging to a server managed via a smart architecture, the specific forwarding configuration applies to the zone on all the servers managed by the smart.
- The forwarding configuration set on a smart architecture zone is automatically inherited by the records it manages. You can override the configuration directly on the physical server zone.
- **Any configuration set at zone level overrides the server or view level configuration.**

To configure forwarding on a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. Open the panel **Forwarding** using **⌵**.
4. Click on **EDIT**. The wizard **Edit a DNS zone** opens.
5. Click on **NEXT** until the page **Forwarding configuration** appears.
6. In the field **Forward mode**, select the mode of your choice according to the table below.

Table 40.1. Forward mode options at zone level

Option	Description
Default	The zone uses the forward configuration set at server or view level.
None	Selecting this option disables the forwarding on the zone.
First	The server sends the queries to the forwarder(s) configured for the zone and, if it does not receive any answer, attempts to find an answer itself.
Only	The server only forwards queries to the forwarder(s) configured for the zone.

7. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displayed the new settings. In the panel **Forwarding**, your configuration is displayed.

You can set a specific forwarding configuration for a zone belonging to a physical server managed via a smart architecture. Keep in mind that:

- When a forward mode is set on a smart architecture, you cannot set the forward mode to *Default* on a zone belonging to a physical server managed via a smart architecture. You can only set a different forward mode.
- Any configuration set at zone level overrides the view or server level configuration.

To configure a specific forward mode on a physical server zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Make sure the servers managed by your smart architectures are displayed. If not, on the right-end side of the menu, click on .
3. At the end of the line of the zone of the physical server of your choice, click on . The properties page opens.
4. Open the panel **Forwarding** using .
5. Click on **EDIT**. The wizard **Edit a DNS zone** opens.
6. Click on **NEXT** until the page **Forwarding configuration** appears.
7. Tick the box **Overwrite the smart settings**. The page refreshes and displays additional fields.
8. In the drop down-list **Forward mode**, you can select *None*, *First* or *Only*. For more details refer to the table [Forward mode options](#).
9. Specify the forwarder(s):
 - a. In the field **Add a forwarder**, type in the address of a forwarder or its name and click on **SEARCH** to retrieve its IP address.
 - b. Click on **ADD** to move it to the list **Forwarders**. Repeat these actions for as many forwarders as needed.
10. Click on **OK** to complete the operation. The report opens and closes. The properties page refreshes and displays the new settings. In the panel **Forwarding**, the Forward value is preceded by the message *Smart configuration is overwritten*.

To revert the specific configuration and inherit it again, edit the *Forwarding* to untick the box *Overwrite the smart settings*.

Configuring DNS Notify Messages at Zone Level

Configuring the Notify at server level allows to set the changes notification once, for all the master zones managed by the view. Once the notification is sent to slave zones, the administrator decides if a zone transfer is relevant. For more details, refer to the sections [Limiting Zone Transfers at Server Level](#) and [Limiting Zone Transfer at View Level](#).

Within SOLIDserver, the notification configuration is done from the panel *Notify* of the properties page. This panel displays:

- The notification type configured for the server,
- The slave zones that should receive the notify messages through their managing server,

- The allow-notify directive configuration of the slave zones. For instance, you can allow all the servers of a network to notify the slave zones of your server or only a few.

Note that there is an implicit allow-notify directive set when you add a slave zone: when you set the Master IP address of the slave zone you are allowing the master zones of this server to send notify messages to your slave zone.

Keep in mind that **any configuration set at zone level overrides the configuration set at server or view level.**

To configure notify messages for a master zone

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the name of the server of your choice. The page **All zones** of the server opens.
3. At the end of the line of the master zone of your choice, click on **ⓘ**. The properties page opens.
4. Open the panel **Notify** using **☰** and click on **[EDIT]**. The wizard opens.
5. If you or your administrator created classes, the **DNS zone class** list is visible. Select a class or *None* and click on **[NEXT]**. The next page of the wizard opens.
6. In the drop-down list **Notify**, configure the zone notification behavior following the table below.

Table 40.2. DNS view notify types

Field	Description
Inherited	The notify messages configuration is inherited from the lowest container for which it was set (view or server). By default, <i>Inherited</i> is selected for each zone.
No	No notify message is sent when changes are performed in the master zones.
Yes	The notify messages are sent to the target of the NS records of the master zone. It is also sent to the IP address(es) specified in the field <i>IP address</i> below.
Explicit	The notify messages are only sent to the IP address(es) specified in the field <i>IP address</i> below.

7. If you selected *Yes* or *Explicit*, you can set the IP address and port of the server(s) which slave zones should receive the messages:
 - a. In the field **IP address**, type in the IP address of another server. The notify message is sent if you chose the notify type *Yes* or *Explicit*.
 - b. In the field **Port**, you can specify the port number that should receive the notify messages on the server you specified in the previous field.
 - c. Click on **[ADD]**. The IP address and port number are displayed in the list **Also notify** as follows: *<ip-address> port: <port-number>*. You can repeat these actions for as many servers as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **[UPDATE]** or click on **[DELETE]** to remove it from the list. If you made changes that you do not want to save, click on **[CANCEL]**.

8. Click on **[OK]** to complete the operation. The report opens and closes. The properties page is visible again. Your notify and also-notify settings are displayed in the panel **Notify**.

To configure notify messages for a slave zone

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Click on the name of the server of your choice. The page **All zones** of the server opens.
3. At the end of the line of the slave zone of your choice, click on **ⓘ**. The properties page opens.
4. Open the panel **Notify** using **☰** and click on **EDIT**. The wizard opens.
5. If you or your administrator created classes, the **DNS zone class** list is visible. Select a class or *None* and click on **NEXT**. The next page of the wizard opens.
6. In the drop-down list **Notify**, set the zone notification type following the table below.

Table 40.3. DNS view notify types

Field	Description
Inherited	The notify messages configuration is inherited from the lowest container for which it was set (view or server). By default, <i>Inherited</i> is selected for each zone.
No	No notify message is sent when changes are performed in the master zones.
Yes	The notify messages are sent to the target of the NS records of the master zone. It is also sent to the IP address(es) specified in the field <i>IP address</i> below.
Explicit	The notify messages are only sent to the IP address(es) specified in the field <i>IP address</i> below.

7. If you selected *Yes* or *Explicit*, you can set the IP address and port of the server(s) which slave zones should receive the messages:
 - a. In the field **IP address**, type in the IP address of another server. The notify message is sent if you chose the notify type *Yes* or *Explicit*.
 - b. In the field **Port**, you can specify the port number that should receive the notify messages on the server you specified in the previous field.
 - c. Click on **ADD**. The IP address and port number are displayed in the list **Also notify** as follows: *<ip-address> port: <port-number>*. You can repeat these actions for as many servers as needed.

You can edit the content of the list. Click on an entry, it is displayed again in the field, edit it and click on **UPDATE** or click on **DELETE** to remove it from the list. If you made changes that you do not want to save, click on **CANCEL**.

8. Click on **NEXT**. The page **Allow notify** opens. It allows to specify if the slave zone can receive master zones notification messages.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 40.4. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.

Type	Description
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL: admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons  and .

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again. Your notify, also-notify and allow-notify settings are displayed in the panel **Notify**.

Managing DNS Security

DNS Security can be configured at zone level through zone transfers configuration, DNS queries restrictions configuration or dynamic update. All these methods set ACLs to allow or deny access to your zones so **keep in mind that the order of the elements listed in the ACL values field is important** as each restriction or permission is reviewed following the order you set in the list.

Dynamic update is detailed in the chapter [Implementing Dynamic Update](#).

Restricting DNS Queries for a Zone

You can specify which hosts are allowed to issue DNS queries for a specific zone. By default, queries are allowed from the local host and the local networks. This property can be configured for an entire server including all zones it contains. For more details regarding restricting DNS queries, refer to the sections [Restricting DNS Queries at Server Level](#) and [Restricting DNS Queries at View Level](#).

Keep in mind that **once the allow query is configured at zone level, it overrides the configuration at server or view level**.

To allow query access for a zone

- In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
- At the end of the line of the zone of your choice, click on . The properties page opens.
- In the panel **Access Control**, click on **EDIT**. The wizard **Allow-query** opens.
- Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 40.5. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

5. Click on **NEXT**. The page **Allow-transfer** appears.
6. Click on **NEXT**. The page **Allow-update** appears.
7. Click on **OK** to complete the operation. The report opens and closes. The parameters are visible in the panel *Access control*, in the list **Allow-query**.

Limiting Zone Transfers for a Zone

DNS zone transfer is a type of DNS transaction employed to replicate and synchronize all copies of the zone used at each server configured to host the zone. SOLIDserver denies zone transfers by default to all DNS server. SOLIDserver supports the allow-transfer zone property that allows to specify which hosts, networks, or TSIG keys are granted or denied the right to do transfers for a specified DNS zone.

Keep in mind that **once the allow-transfer is configured at zone level, it overrides the configuration set at server or view level.**

To allow transfer access for a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Access Control**, click on **EDIT**. The wizard opens: the page **Allow-query** appears.
4. Click on **NEXT**. The page **Allow-transfer** appears.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 40.6. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **▲** and **▼**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

- Click on **NEXT**. The page **Allow-update** appears.
- Click on **OK** to complete the operation. The report opens and closes. The parameters are visible in the panel *Access control*, in the list **Allow-transfer**.

Configuring DNS Update Authorizations on a Zone

Once you authenticated [dynamic update on a server](#), you can complete the configuration via the statement *allow-update* of the master zones. This operation must be done one zone at a time.

Dynamic update replaces or deletes records in a master server by sending it a special form of DNS messages. The format and meaning of these messages is specified in RFC 2136 and indicates which servers or clients are authorized to dynamically update the DNS master zones. By default, all DNS update queries are rejected. Note that:

- You can only complete the master zones configuration from the GUI on EfficientIP and Nominum servers
- You cannot edit the zone update authorizations of Generic servers from the GUI, it has to be done locally.
- You cannot configure dynamic update on Microsoft servers as they do not support TSIG keys. However you can configure them with secure dynamic update using GSS-TSIG keys. For more details, refer to the section [Configuring Secure Dynamic Update](#).

To set up dynamic update, you must:

- Secure the server with a unique TSIG key. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
- Configure the server for dynamic update. For more details, refer to the section [Authenticating the Zones Dynamic Update from the Server](#).
- Configure the statement *allow-update* with the TSIG key:

- If you edited the ACL *admin* of the server, the configuration is complete: by default the ACL *admin* of the physical server is specified in the statement *allow-update* of the master zones.
- If you created an ACL, you must edit the statement *allow-update* and include the new ACL in the permissions.
- If you do not want to rely on ACLs, you must edit the statement *allow-update* and include the new ACL in the permissions.

Keep in mind that allowing updates based on the requestor IP address is insecure, **we strongly recommend using the TSIG key protocol filtering rather than an IP address based filtering.**

To configure the statement *allow-update* of a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **Access Control**, click on **EDIT**. The wizard opens: the page **Allow-query** appears.
4. Click on **NEXT**. The page **Allow-transfer** appears.
5. Click on **NEXT**. The page **Allow-update** appears.

Using the drop-down lists **Type** and **Restriction**, you can configure as many restrictions as you need: grant or deny access to networks, IP addresses, ACLs, and keys. *Type* contains the following options:

Table 40.7. Restrictions and permissions parameters

Type	Description
Network address	Allow or deny an entire network. In the field <i>Network address</i> , type in an IPv4 address following the format <i><ip-address>/<prefix></i> .
IP address	Allow or deny an IP address (of an appliance, user, host...). In the field <i>IP address</i> , type in the IP address.
ACL	Allow or deny an ACL defined at server level in the drop-down list <i>ACL</i> : <i>admin</i> ^a , <i>any</i> , <i>none</i> , <i>localhost</i> and <i>localnets</i> . The ACL list also includes specific ACL created at server level. For more details, refer to the section Configuring Access Control Lists For a Server .
TSIG key	Allow or deny a DNS key defined at server level in the panel <i>Keys</i> .

^aThe ACL *admin* is used by EfficientIP's management platform to configure and exchange data with DNS servers.

Once a restriction/permission is configured as needed, click on **ADD**. The configuration is visible in the list **ACL values**: denied hosts appear preceded by an exclamation mark (!). Make sure the order of the entries in the list suits your needs: the order of the elements listed is important as each restriction or permission is reviewed following the order you set in the list. To organize the list, use the buttons **⬅** and **➡**.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**.

6. Click on **OK** to complete the operation. The report opens and closes. The parameters are visible in the panel *Access control*, in the list **Allow-update**.

To disable dynamic update, refer to the section [Disabling Dynamic Update](#).

Configuring DNS Sources at Zone Level

The Sources configuration is only available for zones managed through an EfficientIP DNS physical server using the SSL protocol. The notify configuration can only be set on master zones, whereas the transfer configuration can only be set on slave zones.

At zone level, configuring DNS sources allows to set physical interfaces that should be systematically used for all notify operations and zone transfer. DNS sources can be inherited from the server and views or set for a zone. The inheritance details are visible the zones properties page.

Keep in mind that **once the sources are configured at zone level, they override the configuration set at server or view level.**

Notify Configuration on Master Zones

From the **Sources** and **Sources V6** panels, through their IP address, you can configure physical interfaces that should be used for the notify options. When editing these panels, you can define the following statements:

notify-source

This statement allows to define the IPv4 address of the physical interface used for all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

notify-source-v6

This statement allows to define the IPv6 address of the physical interface used all the server outgoing notify operations. You can also specify a port for this statement. It is used by master zones and its configuration is therefore displayed on the physical server, views and master zones properties page.

Keep in mind that **once the sources are configured at zone level, they override the configuration set at server or view level.**

To set IPv4 DNS sources on a master zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Type**, type in *Master* to only display master zones.
3. At the end of the line of the zone of your choice, click on **ⓘ**. The properties page opens.
4. Open the panel **Sources** using **⊞** and click on **EDIT**. The wizard **Configuration: Sources** opens.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the notify operations would fail.
 - a. In the field **Notify-source address**, type in the IPv4 address to be used for the outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source port**, you can specify which port on the interface should be used.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

To set IPv6 DNS sources on a master zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Type**, type in *Master* to only display master zones.
3. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
4. Open the panel **Sources** using **⌵** and click on **EDIT**. The wizard **Configuration: Sources** opens.
5. Configure the notify statement. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the notify operations would fail.
 - a. In the field **Notify-source-v6 address**, type in the IPv6 address to be used for the outgoing notify operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Notify-source-v6 port**, you can specify which port on the interface should be used.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

Transfer Configuration on Slave Zones

From the **Sources** and **Sources V6** panels, through their IP address, you can configure physical interfaces that should be used for the transfer options. When editing these panels, you can define the following statements:

transfer-source

This statement allows to determine the IPv4 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

transfer-source-v6

This statement allows to determine the IPv6 address of the physical interface used to execute the zones transfer on the server. You can also specify a port for this statement. It is only valid for slave zones and its configuration is therefore displayed on the physical server, views and slave zones properties page.

use-alt-transfer-source

This statement allows to set the use of an alternate interface IP address for the transfer if the *transfer-source* or the *transfer-source-v6* were to fail. This statement configuration is displayed on the physical server, view and slave zones properties page.

This statement definition is only configurable from the panel **Sources** but applies to interfaces whether they were identified through an IPv4 or an IPv6 address.

Its default value is *no* if the server contains views and *yes* if the server does not contain any view.

alt-transfer-source

This statement allows to determine the alternate IPv4 address of the interface used to execute the zones transfer on the server if the *transfer-source* fails and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

alt-transfer-source-v6

This statement allows to determine the alternate IPv6 address of the interface used to execute the zones transfer on the server if the *transfer-source-v6* failed and if the *use-alt-transfer-source* is enabled. You can also specify a port for this statement. Its configuration is displayed on the physical server, views and slave zones properties page.

Keep in mind that **once the sources are configured at zone level, they override the configuration set at server or view level.**

To set IPv4 DNS sources on a slave zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Type**, type in *Slave* to only display slave zones.
3. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
4. Open the panel **Sources** using **⌵** and click on **EDIT**. The wizard **Configuration: Sources** opens.
5. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the transfer operations would fail.
 - a. In the field **Transfer-source address**, type in the IPv4 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance.
 - b. In the field **Transfer-source port**, you can specify which port on the interface should be used.
 - c. In the drop-down list **Use-alt-transfer-source**, set the use of an alternate interface if need be.

Table 40.8. Use-alt-transfer-source parameters

Parameter	Description
none	This is the default value of the <i>use-alt-transfer-source</i> statement. If your server contains views, it corresponds to <i>no</i> . If your server does not contain any view, it corresponds to <i>yes</i> .
no	This value disables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails.
yes	This value enables the use of an alternate interface if the transfer set via <i>transfer-source</i> or <i>transfer-source-v6</i> fails. In this case, you need to set the alternate interface IP address (and port if you want) through the <i>alt-transfer-source</i> and <i>alt-transfer-source-v6</i> statements in the following steps.

The statement *use-alt-transfer-source* applies to the alternate interfaces declared through IPv4 address (*Alt-transfer-source address*) and IPv6 address (*Alt-transfer-source address-v6*).

- d. If you enabled the use of an alternate interface, in the field **Alt-transfer-source address**, type in the IPv4 address of the alternate interface. It must also be configured on the appliance.
 - e. If you enabled the use of an alternate interface, in the field **Alt-transfer-source port**, you can specify which port on the interface should be used.
6. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

To set IPv6 DNS sources on a slave zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Type**, type in *Slave* to only display slave zones.
3. At the end of the line of the zone of your choice, click on **ⓘ**. The properties page opens.
4. Configure the transfer statements. Make sure to specify the IP address of an interface already declared on SOLIDserver, otherwise all the transfer operations would fail.
 - a. In the field **Transfer-source-v6 address**, type in the IPv6 address to be used for the zones transfer operations. Specify an interface that you already configured on the appliance. If you defined the *use-alt-transfer-source* statement in the panel **Sources**, it applies to the alternate interfaces declared in IPv4 (Alt-transfer-source address) and IPv6 (Alt-transfer-source address-v6).
 - b. In the field **Transfer-source-v6 port**, you can specify which port on the interface should be used.
 - c. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 address**, type in the IPv6 address of the alternate interface. It must also be configured on the appliance.
 - d. If you enabled the *use-alt-transfer-source* in the Sources panel, in the field **Alt-transfer-source-v6 port**, you can specify which port on the interface should be used.
5. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and displays the values you defined.

Chapter 41. Managing DNS Resource Records

The resource records belong to your zones. They usually are contained in master zones and can be replicated to slave zones if need be.

On the page *All RRs*, all records are listed and the column *RR name* indicates if you can edit a record or not.



Figure 41.1. Resource records on the page *All RRs*

- ❶ A record name not underlined indicates that you cannot edit it. In this case, you cannot edit the four records listed because they belong to a physical server managed via a smart architecture.
- ❷ An underlined record name indicates that you can edit it. If you click on a name, you open the wizard *Edit a DNS RR*.

In this case, the underlined records belong to a smart architecture, any changes you make are pushed to the physical server managed via the smart.

- ❸ The SOA record is never underlined because, unlike other records, you cannot edit its properties from the page *All RRs*.
- ❹ This lock indicates that the server records cannot be edited. In this case, the server listed is a physical server managed via a smart architecture.
- ❺ This button, when gray, indicates that you are displaying the records of both the smart architecture and the physical server(s) it manages.

SOLIDserver supports many types of resource records, most can be added to a zone, some are automatically created when you add a zone. All the supported records are listed in the table below.

Table 41.1. Supported DNS resource records

Type	Full name	Description
SOA	Start of Authority	Defines the zone name, an email contact and various time and refresh values applicable to the zone. It is automatically generated when you add a zone and cannot be added manually.
NS	Name Server	Defines the name server(s) that has authority for the domain (defined by the SOA record) or the subdomain. When you add a zone, it is automatically generated once the server has been synchronized
A	IPv4 Address	Maps a host name to an IPv4 address.

Managing DNS Resource Records

Type	Full name	Description
AAAA	IPv6 Address	Maps a host name to an IPv6 address.
AFSDB	AFS Database	Specifies the location of the AFS servers.
CAA	Certificate Authority Authorization	Defines the CAs authorized to issue a certificate for the domain.
CERT	Certificate record	Stores certificates in the DNS.
CNAME	Canonical Name	Defines an alias name for a host.
DHCID	DHCP identifier	Stores DHCP client identifier in the DNS to resolve DNS update conflicts.
DNAME	Delegation of Reverse Names	Sets delegation of reverse addresses primarily in IPv6. (Deprecated, use the CNAME RR instead)
HINFO	System Information	Contains information about a host: hardware type and operating system description.
MINFO	Mailbox mail list Information	Defines the mail administrator for a mail list and optionally a mailbox to receive error messages relating to the mail list.
MX	Mail Exchange	Specifies the mail server/exchanger that services this zone.
NAPTR	Naming Authority Pointer Record	Contains general purpose definition of rule set to be used by applications e.g. VoIP.
NSAP	Network Service Access Point	Defines record (equivalent of an A record) maps a host name to an endpoint address.
OPENPGPKEY	OpenPGP public key record	Publishes the OpenPGP public keys in the DNS.
PTR	Pointer Record	Maps a host name from an IPv4 or IPv6 address, used in reverse mapping (address resolution).
SSHFP	SSH Public Key Fingerprint	Publishes the Secure Shell Fingerprint associated to a specific hostname.
SRV	Service locator	Defines the services available in the zone: LDAP, HTTP, etc...
TLSA	TLSA certificate association	Authenticates TLS client and server entities without a certificate authority.
TXT	Text	Specifies the information associated with a name.
URI	Uniform Resource Identifier	Maps a domain to a URI.
WKS ^a	Well-Known Service	Defines the services and protocols supported by a host. (Deprecated, use the SRV RR instead)
DNSSEC records		
65534	65534	A private type record automatically added to the zone once it is signed with DNSSEC. It cannot be added manually.
CDNSKEY	Child DNSKEY	Contains the public cryptographic key used to sign the zone managing a subdomain with DNSSEC.
CDS	Child DS	Contains information used to verify the validity of the ZSK of a zone managing a subdomain signed with DNSSEC.
DNSKEY	DNS Key	Contains the public cryptographic key used to sign the zone with DNSSEC.
DS	Delegation Signer	Contains information used to verify the validity of the ZSK of a zone signed with DNSSEC.

^aThis record type is obsolete.

When you create a master zone, it automatically contains an SOA record and an NS record. This NS is not generated until the server is synchronized.

Browsing DNS Resource Records

The resource record (RR) is the lowest level of the DNS hierarchy.

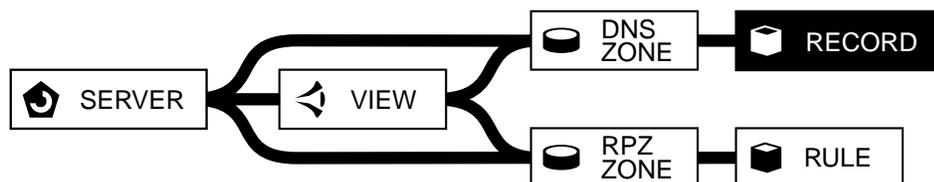


Figure 41.2. The resource record in the DNS hierarchy

If you created RPZ zones, their records; or rules, are listed on the page *All RPZ rules*. For more details, refer to the chapter [DNS Firewall \(RPZ\)](#).

Browsing the DNS Resource Records Database

To display the list of DNS RRs

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. To display the list of DNS RRs of a specific zone, in the column **Zone**, click on the name of the zone of your choice. The page refreshes.

Resource records do not have a properties page, all their information is displayed in the list.

Customizing the Display on the Page All RRs

Users of the group *admin* can create customized column layouts. The button **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the DNS Resource Record Statuses

The column **Status** provides information regarding the resource records you manage.

Table 41.2. DNS resource record statuses

Status	Description
✔ <i>OK</i>	The zone is operational.
⌚ <i>Delayed create</i>	The zone creation is delayed due to a server load or a server unavailability. The creation is automatically pushed when the server is available.
⌚ <i>Delayed delete</i>	The zone deletion is delayed due to a server load or a server unavailability. The deletion is automatically pushed when the server is available.

Adding Resource Records

From the page *All RRs*, you can add almost all the records listed in the table [Supported DNS resource records](#) to a master zone, *Name* or *Reverse*.

When you add master zones, the SOA record and an NS record are created automatically. This NS is only generated after the server is synchronized. You can create other NS records in the zone.

Keep in mind that:

- From one wizard you can create all records. Each record type has a dedicated set of fields to configure.
- All records can be named, if you do not name a record, it is named after the zone it belongs to.
- All the records supported are not supported by all DNS servers:

Table 41.3. Supported record types per server

DNS Server type	Supported records
BIND	A, AAAA, AFSDB, CAA, CDS, CDNSKEY, CERT, CNAME, DHCID, DLV, DNAME, DNSKEY, DS, HINFO, MINFO, MX, NAPTR, NS, NSAP, OPENPGPKEY, PTR, SOA, SSHFP, SRV, TLSA, TXT, URI, WKS.
NSD	A, AAAA, AFSDB, CAA, CNAME, DNAME, HINFO, MINFO, MX, NAPTR, NS, NSAP, OPENPGPKEY, PTR, SOA, SRV, TLSA, TXT, WKS.
AWS Route-53	A, AAAA, CAA, CNAME, MX, NAPTR, NS, PTR, SOA, SPF, SRV, TXT.
Microsoft DNS servers	A, AAAA, CNAME, MX, NS, PTR, SOA, SRV, TXT.

Adding an A Record

The IPv4 Address (A) record maps a host name to an IPv4 address. It can be added to the page *All RRs* of any Master zone. A single host can be mapped toward several A records, or IP addresses, that create an RRset. In this case, the DNS server responds to queries with all the addresses defined but the order depends on the *rrset-order* statement of the server configuration file.

To add an A record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the master zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select **A**.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename*.
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **IP address**, type in the IPv4 Address of the host.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**. The column **RR name** displays its *Complete name*, the column **Value** displays the host IP address you specified.

If you do not name an A record, it takes the same name as the zone it belongs to, which allows DNS clients to find the IPv4 address of your host using only its domain name. This way, querying

the zone name *example.com* would be resolved immediately and provide access to your host through its IP address.

Adding a AAAA Record

The IPv6 Address (AAAA) record, also called *Quad A* record, is used to map a host name to an IPv6 address. It can be added to the page *All RRs* of any Master zone. A single host can be mapped toward several A records, or IP addresses, that create an RRset. In this case, the DNS server responds to queries with all the addresses defined but the order depends on the *rrset-order* statement of the server configuration file.

To add a AAAA record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the master zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select **AAAA**.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **IPv6 address**, type in the IPv6 Address of the host.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**. The **RR name** column displays its *Complete name*, the **Value** column displays the host IP address you specified.

If you do not name an AAAA record, it takes the same name as the zone it belongs to, which allows DNS clients to find the IPv6 address of your host using only its domain name. This way, querying the zone name *example.com* would be resolved immediately and provide access to your host through its IPv6 address.

Adding an AFSDB Record

The Andrew File System Database (AFSDB) record maps a domain name to an AFS database server. Its purpose is to allow to discover the host that provides service AFS within a domain. It is not widely used, an [Adding an SRV Record](#) record could provide the same kind of information.

To add an AFSDB record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select **AFSDB**.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .

6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Sub-type**, type in the version of service AFS used: 1 (AFS version 3.0) or 2 (OSF DCE/NCA version).
8. In the field **AFS server**, type in the AFS hostname.
9. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a CAA Record

The Certificate Authority Authorization (CAA) record allows the holder of a domain to specify which certificate authority(ies) are allowed to issue certificates for the domain.

To add an CAA record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *CAA*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Flags**, type in the number, between *0* and *255*, that influences the interpretation of the record. *0* means that the record is not critical.
8. In the field **Property identifier**, type in the record tag that suits your needs: *issue*, *issuewild*, *iodef* or any other value supported by your DNS engine. For more details, refer to the appendix [DNS Resource Records Related Fields](#).
9. In the field **Value**, type in the domain of the CA associated with the tag or a URL, depending on the tag specified in the *Property identifier*.
10. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a CERT Record

The Certificate (CERT) record allows to store certificates in the DNS.

To add an CERT record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *CERT*.

5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Type** type in the number, between *0* and *65535*, that specifies the type of certificate. For more details, refer to the appendix [DNS Resource Records Related Fields](#).
8. In the field **Key tag**, type in the certificate's key tag, a 16-bit value computed for the key embedded in the certificate.
9. In the field **Algorithm**, type in the public key's cryptographic algorithm.
10. In the field **Certificate or CRL**, type in the certificate, encoded in base-64 format.
11. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a CNAME Record

The Canonical Name (CNAME) record maps an alias to a real name, also called *canonical name*. This name may lie inside or outside your zone but it generally exists elsewhere in your DNS. The CNAME is mostly used if a host has several possible names, the alias provides a way of saving all the possible names in your zone to resolve more easily IP or domain name queries. The CNAME always points to another record, usually an A record. During a query, the CNAME returns the canonical name and IP address embedded in the A record. That's why a CNAME should not point to another CNAME record, the DNS answer would take longer and could overload the server: the first CNAME would point to another CNAME that would point to another CNAME and so forth until finally getting the IP address from the A record.

Keep in mind that **each CNAME RR name is unique**: you cannot have several records named *www* in the same zone. Their complete name would be *www.example.com* and as the CNAME is an alias, it should provide a link toward a canonical name that has not been declared in the zone yet.

To add a CNAME record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the master zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *CNAME*.
5. In the field **RR name**, name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Hostname**, type in the host of your choice canonical name.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

There should be as many hostname aliases as there are CNAME records in your zone.

Adding a DHCID Record

The DHCP identifier (DHCID) record allows to store DHCP client identifier (DUID) in the DNS to resolve DNS update conflict performed by DHCP clients.

To add an DHCID record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *DHCID*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Key**, type in the encoded DUID in a binary format, encoded in base-64
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**.

Adding a DNAME Record

The Delegation of Reverse Names (DNAME) record is used to map DNS subdomains with each other. It does not redirect a query toward a subdomain: it rewrites the query. Technically, it rewrites the subdomain query suffix and looks for a record within the zone matching this new name. It is especially useful if a company has changed domain name or reorganizes its subdomains management.

Keep in mind that a DNAME record rewrites the subdomain suffix and applies to all its subdomains. A DNAME record rewriting a query from *support.company.com* to *support.company.corp* also applies to queries for *fr.support.company.com* or *es.support.company.com* . The DNAME configuration applies to any label located left of the specified domain name.

A zone configured with a DNAME has records that send back the proper information to DNS clients. If the value of the DNAME is *support.company.corp*, there should be an A record, for instance, named *support.company.corp* providing an IP address clients can reach.

Keep in mind that unlike a CNAME, the DNAME points a name and not to a record within the zone.

To add a DNAME record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *DNAME*.

5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Domain**, type in the domain name of a subdomain of the zone.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a DNSKEY Record

The Domain Name System KEY (DNSKEY) record is used in zones signed with DNSSEC and contains the public cryptographic key (KSK or ZSK) used to validate signatures.

You do not need to add a DNSKEY record if you signed zones from the GUI, it is created automatically for each zone.

If you manage from the GUI an external DNS server that contains one or more zones already signed with DNSSEC, you can add a DNSKEY record to each concerned zone. When the zone signature is not performed using the appliance, SOLIDserver cannot push the DNSSEC keys to the server and displays them like any other signed zone. Therefore:

- The DNSKEY record is not listed among the records of the zones.
- The DNSSEC keys of each zone are not listed on the page *All DNSSEC keys*.
- The zones are not displayed as DNSSEC compliant even though they are.

Adding a DNSKEY record can therefore ease up the zone management and ensure that the zones you signed from another platform are marked *yes* in the column *DNSSEC*. For more details, refer to the chapter [DNSSEC](#).

To successfully add a DNSKEY record, you need the DNSKEY flags, protocol, algorithm and key; all of which are available in txt file generated after the zone signature.

To add a DNSKEY record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *DNSKEY*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Flags**, type in or paste the zone key flag.
8. In the field **Protocol**, type in or paste the protocol value.
9. In the field **Algorithm**, type in or paste the public key's cryptographic algorithm.

10. In the field **Key**, type in or paste the public key material.
11. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a DS Record

The Delegation Signer (DS) record is a DNSSEC that creates the chain of trust or authority from a signed parent to a child zone. It can be used to verify the validity of the ZSK of a subzone. It is composed of the parent zone key tag, key algorithm, digest type and digest itself. For more details, refer to the section [Managing DNSSEC on Authoritative Servers](#).

To add the DS records, refer to the section [Publishing the Delegation Signer in the Parent Zone](#).

Adding an HINFO Record

The System Information (HINFO) record allows to specify the server type of CPU and OS in use. This record information can be used by some application protocols (like FTP). This record is rarely used on public servers.

Keep in mind that if you name an HINFO record like an A or AAAA record, they are linked together in the zone file and provide additional information when the domain name they share (an identical Complete name in the GUI) is queried.

To add an HINFO record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *HINFO*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. Next to the field **CPU**, select a CPU in the drop-down list. It is displayed in the field. If your CPU is not listed, select *Other* in the drop-down list and type in your CPU in the field.
8. Next to the field **OS**, select an OS in the drop-down list. If your OS is not listed, select *Other* in the drop-down list and type in your OS in the field.
9. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

The HINFO can also be used as a specific TXT record and contain other information.

Adding an MINFO Record

The Mailbox mail list Information (MINFO) record defines the mailbox administrator for a mail list or even the mailbox that should receive error messages relating to the mail list.

To add an MINFO record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *MINFO*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Responsible email**, type in the email address of the administrator of the mail list.
8. In the field **Error email**, type in the email address of the person that should receive the error messages regarding the mail list.
9. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**.

Adding an MX Record

The Mail Exchanger (MX) record allows to set the name and relative preference of your mail exchangers, in other words, mail servers for the zone. Keep in mind that:

- **An MX record should not point to a CNAME record**, as detailed in section 10.3 of RFC 2181. Therefore, if you have a CNAME called *mail* for the zone *example.com* (its complete name would be *mail.example.com*) and if one of your mail exchangers name is *mail.example.com*, you need to remove the alias from the zone to be able to declare the mail exchanger name in the MX record. To make the answer for the MX more efficient, you should also add an A or AAAA record pointing to the IP address of the mail server.
- **If the mail server stated in one of the MX records lies in the zone, you should add an A record**. This A record name becomes the mail server and its value becomes its IP address.

To add an MX record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *MX*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Preference**, type in a number between 0 and 65535, which defines which server has priority if there are several MX records in the zone. The lowest the value has the priority over the other server(s), it can be 0.

8. In the field **Mail server**, type in the mail server hostname.
9. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

You can add as many MX records as you need in your master zones, it all depends on the number of mail exchangers you want to declare.

Adding an NAPTR Record

The Naming Authority Pointer (NAPTR) record is a Dynamic Delegation Discovery System (DDDS) record used to define a rule that may be applied to private data owned by a client application, as detailed in RFC 3403.

To add an NAPTR record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *NAPTR*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour **. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Order**, type in a number between 0 and 65535 to define which RR has priority if there are several NAPTR records in the zone. The lowest value has the priority over the other record(s).
8. In the field **Preference**, type in a number between 0 and 65535 to define which RR has priority if several NAPTR records have the same *Order* in the zone. The lowest value has the priority over the other record(s).
9. In the field **Flags**, type in the string that corresponds to the action you want your client application to perform.
10. In the field **Service**, type in the services parameters needed according to your client application syntax.
11. In the field **Regex**, type in the string that contains a substitution expression to be applied to the original string specified in the field *Flags*.
12. In the field **Replace**, type in the FQDN domain name to be queried when looking for the potential data specified in the field *Flags*.
13. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding an NS Record

The Name Server (NS) record is used to list all the DNS name servers that have authority over a zone. NS records must be declared both in the parent and the child zones. In the parent zone, they indicate the zone authoritative server, in the child zone where they constitute the point of delegation.

The requirement is that at least two name servers are defined for each public domain, so there is at least two NS records in each zone. The first NS record, named after the zone is created automatically when you create zones from the GUI to indicate the authoritative server; all other NS records must be added manually following the procedure below.

We strongly recommend that you create an A record for each NS server to provide detailed information to the domain name query. This process is called creating a *glue record*, that way once your domain is queried, it can return its authoritative servers name and IP address.

Keep in mind that RFC 2181 stipulates that the **NS record can point to other records but never to a CNAME record** as the query answer does not return an address with the NS record and in some cases might make the query fail altogether.

To add an NS record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *NS*.
5. In the field **RR name**, name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **DNS server**, type in the DNS server hostname.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is .

Adding an NSAP Record

The Network Service Access Point (NSAP) record maps a hostname to an endpoint address for the ISO's Open Systems Interconnect (OSI) system, in that sense it is the equivalent of an A record. The NSAP RR is described in the RFC 1706.

To add an NSAP record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *NSAP*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.

7. In the field **Name**, type in the NSAP address of the end system. It should start with *0x* and not exceed 255 hexadecimal characters separated by dots.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding an OPENPGPKEY Record

The OpenPGP public key record (OPENPGPKEY) record allows to publish OpenPGP public keys in the DNS for a specific email address.

To add an OPENPGPKEY record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *OPENPGPKEY*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour**. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Key**, type in the OpenPGP public key, without the last line.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a PTR Record

The Pointer (PTR) record is used to reverse map an IP address to a host name and can be used both in IPv4 and IPv6. These records can only be added to reverse zones, they basically provide the exact opposite information than the A and AAAA records.

The PTR name is always displayed in the RR name column in reverse with the syntax *B4.B3.B2.B1.in-addr.arpa* but it is treated like a name. Which is why it is possible to set IP addresses final section (B4) with a value that does not respect the IP protocol: a value greater than 255 in IPv4 and greater than *ffff* in IPv6. This lack of limitation in the interface provides an additional tool for specific configurations.

The PTR being used for reverse host name look ups, it does not make sense to name multiple PTR records with the same name, i.e. same IP address. However, to provide reverse round-robin configuration, you can set several IP addresses with different values. For more details, refer to the section [Load Balancing with Round-Robin](#).

To add a PTR record in a reverse zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the master *reverse* zone of your choice to display the RRs it contains.
3. In the menu, click on  **Add > RR**. The wizard **Add a DNS RR** opens.

4. In the drop-down list **RR type**, select *PTR*.
5. Set the IP address in reverse via the field *RR name* or the field *IP address*. You must fill in one of the two fields:
 - a. If you want to use the field **RR name**, you can type a number corresponding to the remaining section of the IP address of your choice. Filling in this field empties the field *IP address* as only one of the two is required. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.reversezonename* .
 - b. If you want to use the field **IP address**, the first sections of the IP address that you set for the reverse zone is displayed, note that is not displayed in reverse. Type in the missing dot and final section of the IP address. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.reversezonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Hostname**, type in the hostname that should be returned when the IP address you stated above is queried.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

All the addresses used to name your PTR records provide as many entries toward the host names of your choice in your reverse master zones.

Adding an SSHFP Record

The SSH Public Key Fingerprint (SSHFP) record allows to publish the Secure Shell (key) Fingerprint associated to a specific host. That host is specified in the RR name.

To add an SSHFP record to a zone

1. In the sidebar, go to  **DNS** > **RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *SSHFP*.
5. In the field **RR name**, specify the name of the host associated with the secure shell fingerprint of the record. The field **Complete name** auto-completes and displays the RR full name as follows: *<host>.<domain>* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Algorithm**, type in the number matching the algorithm used by the Public Key: 1 (RSA), 2 (DSA), 3 (ECDSA) or 4 (Ed25519).
8. In the field **Type**, type in the number matching the type of fingerprint used: 1 (SHA-1) or 2 (SHA-256).
9. In the field **Fingerprint**, type in the hexadecimal string of the hash result.

- Click on to complete the operation. The report opens and closes. The record is now listed and marked  **OK**.

Adding an SRV Record

The Services (SRV) record allows to associate a service with a hostname. That way, users can locate a service via the relevant SRV record. The answer to a successful SRV query should provide the user with a hostname, the port providing the service and the hostname priority. If there are several hosts in the zone, their weight defines which one should be used..

This record only allows one piece of information per field, so if for instance you want to configure a set of ports for one service, you can create several SRV records each with the same information in all fields except the port, priority and weight.

To add an SRV record to a zone

- In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
- In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
- In the menu, click on  **Add**. The wizard **Add a DNS RR** opens.
- In the drop-down list **RR type**, select **SRV**.
- In the field **RR name**, name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
- In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
- In the field **Priority**, type in a number, between 0 and 65535, to define which server has priority if there are several SRV RRs in the zone. The lowest the value has the priority over the other server(s).
- In the field **Weight**, type in a number, between 0 and 65535, to define the server weight. If two SRV RRs have the same priority, the weight defines which server should be more used. The greater the value is, the more the server is solicited. Basically, it gives priority to the SRV RR with the greatest weight value. If you type in 0, there is no weighting.
- In the field **Port**, type in the port number that delivers the service to the target.
- In the field **Target**, type in the hostname of the server delivering the service.
- Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a TLSA Record

The TLSA certificate association (TLSA) record allows to authenticate TLS client and server entities without a certificate authority (CA). It is used to associate a TLS server certificate or public key with the domain name where the record is found.

To add an TLSA record to a zone

- In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
- In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.

3. In the menu, click on **+** **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *TLSA*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Certificate Usage**, type in a number, between *0* and *255*, that identifies the association data provided to match the certificate presented in the TLS handshake. For more details, refer to the appendix [DNS Resource Records Related Fields](#).
8. In the field **Selector**, specify the part of the TLS certificate that the server presents to compare with the association data: *0* (the Full certificate) or *1* (the SubjectPublicKeyInfo).
9. In the field **Matching Type**, specify the way the association data is presented, where *0* is the Exact match on selected content, *1* is an SHA-256 hash of selected content and *2* is an SHA-512 hash of selected content.
10. In the field **Certificate Association Date**, specify the *certificate association data* to be matched. The expected value in the field is a string of hexadecimal characters, it can include spaces.
11. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a TXT Record

The Text (TXT) record allows to associate text with a name in your zone. You can use the TXT record value to describe a host, provide services contacts or even define the Sender Policy Framework (SPF) information record.

To add a TXT record to a zone

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+** **Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *TXT*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Text**, type in the text of your choice. This field text can contain a maximum of 2100 characters, including spaces.
8. Click on to complete the operation. The report opens and closes. The record is now listed and its status is  **OK**.

Adding a URI Record

The Uniform Resource Identifier (URI) record allows to map a domain to a URI.

To add an URI record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *URI*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **Priority**, type in a number, between *0* and *65535*, to define which target URI has priority if there are several URI records in the zone. The lowest value has the priority over the other URI.
8. In the field **Weight**, type in a number, between *0* and *65535*, to define the target URI weight. If two URI records have the same priority, the weight defines which one is more used. The field gives priority to the URI record with the greatest weight value: the greater the value is, the more the URI is solicited. If you type in *0*, there is no weighting.
9. In the field **Target**, type in targeted URI following the format described in RFC 3986.
10. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**.

Adding a WKS Record

The Well-Known Services (WKS) record is used to define the services and protocols used by a host.

This record type is not **obsolete**, the **SRV** record can provide the same information.

To add a WKS record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select *WKS*.
5. In the field **RR name**, you can name your RR. The field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **IP address**, type in the IPv4 Address of the host that contains the services listed in the field *Services*.
8. In the drop-down list **Protocol**, select the protocol that suits your needs.
9. In the drop-down list **Services**, select the service that suits your needs.

- Click on **OK** to complete the operation. The report opens and closes. The record is now listed and its status is **OK**.

Editing Resource Records

You can edit all the records contained in a master zone.

From the list *All RRs*, you can differentiate the ones you can or cannot edit, refer to the image [A list of resource records](#) above. For instance, the records of a physical server managed via a smart architecture cannot be edited.

Note that SOA records cannot be edited like the other records: you can tick them to change their TTL or value, but you cannot edit them individually from the page. For more details, refer to the sections [Editing Several Records at Once](#) and [Editing the SOA from the Zone Properties Page](#).

Editing One record

From the page *All RRs*, you can edit most records by clicking on their *Name*. Note that:

- You cannot edit the records of a physical server managed by a smart architecture. You need to edit the record of the smart, all your changes are then pushed on the physical server(s) managed by the smart architecture
- The SOA records cannot be edited that way. To edit an SOA record, refer to the sections [Editing Several Records at Once](#) and [Editing the SOA from the Zone Properties Page](#).

To edit a resource record

- In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
- In the column **Zone**, click on the name of the zone of your choice to display the records it contains.
- In the column **RR name**, click on the name of the RR of your choice. The wizard **Edit a DNS RR** opens.
- Edit, if need be, the values and **TTL** of the record following the table appropriate procedure in the [Adding Resource Records](#) section above. The default TTL of an RR is *1 hour*.

Keep in mind that if several records in a zone share the same name, editing the TTL on one also edits the TTL on the records sharing that name.

- Click on **OK** to complete the operation. The report opens and closes. The changes are visible on the page.

Editing Several Records at Once

From the page *All RRs*, you can tick any record, SOA included, to edit its *TTL* or *Value*. These changes can also be performed on several records at once.

To replace the TTL of one or several resource records

- In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
- In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
- Tick the record(s) for which you want to replace the TTL.

4. In the menu, select  **Edit > Replace > The TTL of an RR**. The wizard **Replace the TTL of an RR** opens.
5. In the field **TTL**, specify the expiration time of the record in seconds or use one of the pre-defined values in the drop-down list. The default TTL for an RR is *1 hour*.
6. Click on to complete the operation. The report opens and closes. The page refreshes, the new TTL is visible.

You can tick several records and edit their Value. The wizard allows to replace existing values, you can specify whole values or part of a value, for instance a domain name stated in all the records. Note that the wizard returns an error if you specify a value that does not exist.

To replace the value of one or several resource records

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. Tick the record(s) for which you want to replace a value.
4. In the menu, select  **Edit > Replace > The value of an RR**. The wizard **Replace the value of an RR** opens.
5. In the field **Replace**, type in the value you want to replace.
6. In the field **By**, type in the new value that should replace the content of the field *Replace*.
7. In the drop-down list **Exact search**, you can either select *Yes* or *No*, refer to the table below.

Table 41.4. DNS resource records replacement options

Type	Description
Yes	Select this option if the value specified in the field <i>Replace</i> must be replaced as a whole. Keep in mind that it is then considered as unique, so if the RR contains the same value several times, each of them has to be replaced individually. By default, <i>Yes</i> is selected.
No	Select this option if every occurrence of the value specified in the field <i>Replace</i> must be replaced every time it appears in the column Value of the RR.

8. Click on to complete the operation. The report opens and closes. The page refreshes, the changes are visible in the list.

Editing the SOA from the Zone Properties Page

As the SOA is automatically generated by SOLIDserver when you add a Master zone, you can edit it from the properties page of the zone. Keep in mind that the SOA contains information critical to the zone and editing it can have heavy consequences on a zone management.

The SOA contains the zone's serial number, administrator email and configuration information (refresh, retry, expiration...) all of which you can edit. Note that:

- You cannot rename an SOA, it is automatically named like the zone when you create it.
- The SOA also contains the name of the primary server, the DNS server that has authority over the zone, but you cannot edit it.

To edit the *TTL* or *Value* of one or several SOA records, refer to the section [Editing Several Records at Once](#).

To edit an SOA record

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the zone of your choice, click on **⌵**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS zone** opens.
4. Click on **NEXT** until you get to the last page of the wizard.
5. Edit the SOA parameters, according to your needs:

Table 41.5. DNS zone advanced parameters

Field	Description
Primary server	Specify the primary Master server for the zone. When you create a zone on a smart server, it is automatically filled and cannot be edited.
Responsible	Specify the administrator email address for the zone.
Serial number	The zone serial number. It is automatically incremented for each zone change.
Refresh	The refresh period for slave server(s). When the selected period is reached, the slave server(s) read the Master SOA record and, if their data is not up-to-date, trigger a zone transfer to get the latest version of the zone.
Retry	The retry interval if the server fails to reach the master during a refresh cycle.
Expiration	The period after which the records are considered to be no longer valid/authoritative and the server stops responding to queries for the zone.
Minimum	The negative caching period of the zone, in seconds. This period is used as TTL for every NXDOMAIN returned to clients querying unexisting records.
TTL	The TTL (Time to Live) of the SOA, its duration, in seconds. The drop-down list next to the input field allows to select durations in human-readable format.

6. Click on **OK** to complete the operation. The wizard closes. The page refreshes, the changes are listed.

Configuring the Delegation at Record Level

At RR level, the delegation parameters are managed through the Start of Authority (SOA), the Name Server (NS) and the Address (A or AAAA) records. The SOA and NS records are generated upon creation of a zone.

Note that the primary NS record of a zone is generated once the server is synchronized and indicates the authoritative server of the zone.

Delegating a sub-domain simply consists of adding both an NS and an A (or AAAA) RR in the parent zone pointing to the sub-domain:

- The NS record indicates which servers are authoritative for the zone. You can also create additional NS records to delegate authority for the zone to other DNS servers.
- The A / AAAA record indicates the IP address of the server that has authority over the sub-domain and therefore needs to be added in the RRs list of the parent zone.

Let's consider the zones *efficientip.com* and *support.efficientip.com* for the purpose of illustrating the delegation configuration. The parent zone, *efficientip.com*, is managed through the server *ns1.efficientip.com* and the child zone, *support.efficientip.com*, is managed through *ns2.efficientip.com*. You need to add the relevant records in the parent zone. On the one hand, add the NS record, name it *support* (it is then listed as *support.efficientip.com* as the RR name auto-completes

with the domain name at the end) and indicate the server that has authority over it in the adequate field, in our case *ns2.efficientip.com*. On the other hand, add the A record named *ns2* (once again, its name auto-completes with the zone name and obtain the server actual name) and indicate its IP address. That way, you should have two new records in the parent zone: an NS RR, *support.efficientip.com*, pointing toward the delegated child zone and a glue A record, *ns2.efficientip.com*.

To add an NS record to a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Name**, click on the name of the zone of your choice. The page **All RRs** of the zone opens.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the field **RR type**, *NS* is displayed.
5. In the field **RR name**, name your RR after the sub-domain. Note that the **Complete name** field auto-completes and displays the RR full name as follows: *RRname.zonename* .
6. In the field **DNS server**, type in the name of the server that has authority over the sub-domain.
7. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and marked **OK**.

To add an A record to a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the column **Name**, click on the name of the zone of your choice. The page **All RRs** of the zone opens.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the field **RR type**, *A* is displayed.
5. In the field **RR name**, name your RR after the server that has authority over the sub-domain (the same one as the DNS server specified when adding the NS record). Note that the field **Complete name** auto-completes and displays the RR full name as follows: *RRname.zonename* and should match the server name.
6. In the field **TTL**, specify an expiration time of the record in seconds. The default TTL for an RR is *1 hour* *. You can edit it if need be using the field on the left or one of the values listed in the drop-down list on the right.
7. In the field **IP address**, type in the IP address of the authoritative server.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed and marked **OK**.

Copying or Moving Records

At some point you can use the option *Migrate* to move or copy RR(s) from one DNS server or view to the other.

Note that this option has nothing to do with the zones database replication of the DNS command *allow-transfer*. Duplication and migration of a zone includes the records it manages.

To copy a record

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. Tick the record(s) you want to duplicate.
3. In the menu, select **Edit > Migrate**. The wizard **Copying/Moving RRs** opens.
4. In the drop-down list **Method**, select *Copy*.
5. In the drop-down list **Target server**, select the server of your choice. The **Target zone** drop-down list appears.
6. In the drop-down list **Target zone**, select the zone of your choice. If you created views in your server, the zone is named *zone (view)*.
7. In the drop-down list **Existing records**, choose if you want to overwrite RRs with the same name. The wizard refreshes.
8. Click on **OK** to complete the operation. The report opens and closes. The page **All RRs** is visible again and displays the migrated record. Note that the complete name of the RR(s) in the column **RR name** is now *RRname.newzonename*.

To move a record

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. Tick the record(s) you want to move.
3. In the menu, select **Edit > Migrate**. The wizard **Copying/Moving RRs** opens.
4. In the drop-down list **Method**, select *Move*.
5. In the drop-down list **Target server**, select the server of your choice. The **Target zone** drop-down list appears.
6. In the drop-down list **Target zone**, select the zone of your choice. If you created views in your server, the zone is named *zone (view)*.
7. In the drop-down list **Existing records**, choose if you want to overwrite RRs with the same name.
8. Click on **OK** to complete the operation. The report opens and closes. The page **All RRs** is visible again and displays the migrated record. Note that the complete name of the RR(s) in the column **RR name** is now *RRname.newzonename*.

Changing the Hostname Convention

At any time, you can change the RR naming convention to allow or prohibit the use of some characters or patterns in the records full name. Before editing the record naming convention, keep in mind that the hostname naming convention stated in the RFC 1034 only allows alphanumeric characters and hyphens. It does not include other special characters, such as underscore ("_"). Therefore, dynamic updates from Microsoft Active Directory controllers might not be accepted.

The naming convention can be set from the GUI via one global regular expression (or regex) that applies to the name of all the DNS zones resource records. **SOLIDserver default naming convention regular expression** allows all characters including the hyphen ("-"), the dot (".") and the underscore ("_"):

```
(^[*][.])?[-_a-z0-9\u00c0-\uffff]+([.][-_a-z0-9\u00c0-\uffff]+)*$|(^[*]$)|(^$)
```

You can change this regular expression from the registry database following the procedure below.

To change the records naming convention

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in the keyword *rr_name*. The list is filtered and displays the item *www.display.checktype.regex.js.rr_name*.
4. In the column **Value** of the entry *www.display.checktype.regex.js.rr_name*, click on the *<regular_expression>*. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, edit the regex to match your needs. [The regex default value is detailed above](#).
6. Click on to complete the operation. The report opens and closes. The **Registry database** is visible again.

Deleting Resource Records

Except for the basic SOA and NS records generated during the creation of a zone, all the resource records that you created within a zone can be deleted.

To delete a resource record

1. In the sidebar, go to  **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display the RRs it contains.
3. Tick the record(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The RR is marked  **Delayed delete** and is then no longer listed.

Load Balancing with Round-Robin

The load balancing or Round-Robin functionality is useful if you have a number of equivalent network resources, like mirrored FTP servers, Web servers, and would like to spread the load among them. You establish one domain name that refers to the group of resources, configure clients to access that domain name, and the name server inverse-multiplexes the accesses between the IP addresses you list.

For example, if you have three *www* servers with network addresses of 10.0.0.1, 10.0.0.2 and 10.0.0.3, a set of A resource records means that clients connect to each machine one third of the time. When a resolver queries for these records, BIND rotates them and respond to the query with the records in a different order. In the example above, clients randomly receive records in the order 1, 2, 3; 2, 3, 1; and 3, 1, 2. Once the query is answered a first time with 1, the next client querying the same name receives a different answer: 2; and so forth. There is no configuration needed, the balancing is automatically activated if three different servers resolve to the same domain name (to follow the example: *www.yourdomain.com*).

Note that if you configured the DNS advanced properties *Create PTR*, a PTR record is created only for the first A record you add.

SPF Record

Sender Policy Framework (SPF) is an email validation system designed to prevent email spam by detecting email spoofing using the senders' IP address. SPF allows administrators to specify which hosts are authorized, or not, to send emails from a given domain. Mail exchangers use the DNS to verify that the host sending the email from a given domain is sanctioned by that domain's administrators. The SPF record is actually a single string of text found in the value of a single TXT record.

In 2003, when SPF was first being developed, the requirements for assignment of a new DNS RR type were considerably more stringent than they are now. In its review of the RFC 4408, the IETF SPFbis working group concluded that its dual RR type transition model was fundamentally flawed since it contained no common RR type that implementers were required to serve and required to check.

The Simple Mail Transfer Protocol allows any computer to send email claiming to be from any source address. This is exploited by spammers who often use forged email addresses, making it more difficult to trace a message back to its sender, and easy for spammers to hide their identity in order to avoid responsibility. It is also used in phishing techniques, where users can be duped into disclosing private information in response to an email purportedly sent by an organization such as a bank. SPF allows the owner of an Internet domain to specify which computers are authorized to send mail with sender addresses in that domain. Receivers verifying the SPF records may reject messages from unauthorized sources before receiving the body of the message. The sender address is transmitted at the beginning of the SMTP dialog. If the server rejects the sender, the unauthorized client should receive a rejection message, and if that client was a relaying message transfer agent (MTA), a bounce message to the original sending address may be generated. If the server accepts the sender, and subsequently also accepts the recipients and the body of the message, it should insert a field Return-Path in the message header in order to save the sender address. While the address in the Return-Path often matches other originator addresses in the mail header such as From or Sender, this is not necessarily the case, and SPF does not prevent forgery of these other addresses.

Examples of SPF records

```
example.com. IN TXT "v=spf1 ip4:192.0.2.0/24 ip4:198.51.100.123 a -all"
```

```
example.com. 21600 IN SPF "v=spf1 +all"
```

```
IN TXT "v=spf1 mx -all"
IN TXT "v=spf1 redirect=_spf.example.com"
IN TXT "v=spf1 ip4:192.0.2.0/24 ip4:198.51.100.123 a -all"
IN TXT "v=spf1 include:example.org -all"
```

Chapter 42. Implementing Dynamic Update

Dynamic Domain Name Server (DDNS) is the system, detailed in RFC2136, through which IP address assignment updates in the DHCP are immediately reflected in the DNS records for the hosts. DDNS enables a DNS server to accept and receive an update every time a dynamic client changes their IP address. Updating the DNS dynamically eliminates the need for an administrator to manually set large numbers of records as authorized users can add, delete, and edit records on the fly.

In the wrong hands, dynamic updates can open up your network to certain vulnerabilities as users could update some or many of the records on a DNS server organization with incorrect information. Which is why, within SOLIDserver, DDNS relies on the Transaction SIGnature (TSIG). Described in RFC 2845, TSIG is based on the use of a symmetrical key, and we create one update key per Efficient IP DNS server. That way, every transaction from DHCP servers to DNS servers is automatically protected.

SOLIDserver supports both dynamic update via TSIG keys and secure dynamic update via Generic Security Service Algorithm for Secret Key Transaction (GSS-TSIG) for DNS servers serving Microsoft Active Directory clients.

You can configure your EfficientIP DNS server to authenticate your AD clients and grant or deny them dynamic update rights on the master zone managing your AD server domain via the statement `allow-update-policy`. For more details, refer to the section [Configuring Secure Dynamic Update](#).

Configuring Dynamic Update

To set dynamic update, you must configure your DNS server and then edit the statement `allow-update` of master zones of your choice. The configuration requires to:

1. Secure the server. You must edit the server main properties and specify a TSIG key. For more details, refer to the section [Securing the Management of DNS Servers Using a TSIG Key](#).
2. Configure the server for dynamic update.
 - Either edit the default ACL `admin` to include the same TSIG key to its permissions.
 - Or create a new ACL that includes the TSIG key to its permissions. For more details, refer to the section [Configuring Access Control Lists For a Server](#).

For more details, refer to the section [Authenticating the Zones Dynamic Update from the Server](#).

3. Configure the statement `allow-update` of your master zones with the same TSIG key. For more details, refer to the section [Configuring DNS Update Authorizations on a Zone](#).

Note that if you edited the ACL `admin` of the server, the configuration is complete because, by default, the ACL `admin` of the physical server is specified in the statement `allow-update` of the master zones.

You can configure dynamic update on all DNS servers except Amazon Route 53 DNS servers.

Configuring Secure Dynamic Update

Secure dynamic update for AD domains can be configured thanks to GSS-TSIG keys. You can configure an EfficientIP DNS server managing your AD domain to authenticate AD users. Once the server is set, you can configure the statement *update-policy* on the master zone managing your AD domain to allow or deny dynamic update of your DNS server for AD users.

To properly set up secure dynamic update you must:

1. Meet the prerequisites.
2. Make sure the master zone managing your AD domain is properly configured.
3. Create and upload the GSS-TSIG key.
4. Configure your server with the GSS-TSIG key.
5. Configure the statement *update-policy* on the zone managing your AD domain.

Prerequisites

- SOLIDserver version 6.0.0.
- A Windows Active Directory server:
 - Without DNS server, your DNS server is an EfficientIP DNS server.
 - With at least one user with sufficient DNS rights, able to create GSS-TSIG keys. This user must share the same name as your EfficientIP DNS server following the format *my-ad-user.my-ad-domain*.
 - With DES certificate enabled. For Windows Servers 2008 R2, it is disabled by default so you must enable *DES_CBC_MD5*, *AES128_HMAC_SHA1* and *AES256_HMAC_SHA1* in the *Security Options* of the *Local Security Policy*.
- At least one EfficientIP DNS server named after an AD user with sufficient DNS rights. The server is named following the format *my-ad-user.my-ad-domain*.

The name of the DNS server can be set based on an existing user, or an AD user can be named after the DNS server name. Either way, they must have the same name and that name must include the full AD domain. To create a server, refer to the procedure [To add an EfficientIP DNS server](#).

- The EfficientIP DNS server must contain a master zone managing your AD domain, so it named with a format similar to *my-ad-domain.corp*. This master zone must be configured with 9 key records (1 SOA, 2 A and 6 SRV) linking your AD server, AD domain and DNS server. To create the zone, refer to the procedure [To add a master zone](#).
- SOLIDserver and your AD server must be set at the same time. You should configure the same NTP server on SOLIDserver and your AD server.

Limitations

- Secure dynamic update can only configured for EfficientIP DNS and EfficientIP DNS Package servers.
- For each physical DNS server there must be an AD user. Therefore, to configure secure dynamic update for a Multi-Master smart architecture, you need two AD users to generate the GSS-TSIG key for the server named after them.

Making Sure your Master Zone is Properly Configured

As detailed in the [Prerequisites](#), the master zone managing the AD domain has to be configured with records linking the domain, the AD server and the DNS server. Without these records, you cannot set up the AD authentication for AD clients' dynamic update.

The zone configuration must contain the 9 records configured as detailed in the example below.

```
ad-domain.corp SOA TTL dns-server.ad-domain.corp X X X X X X X dns-server.ad-domain.corp A TTL IP
ad-server.ad-domain.corp A TTL IP _kerberos._udp.ad-domain.corp SRV TTL 0 100 464
ad-server.ad-domain.corp _kerberos._udp.ad-domain.corp SRV TTL 0 100 88 ad-server.ad-domain.corp
_kerberos._tcp.ad-domain.corp SRV TTL 0 100 464 ad-server.ad-domain.corp _kerberos._tcp.ad-domain.corp
SRV TTL 0 100 88 ad-server.ad-domain.corp _ldap._tcp.pdc._msdcs.ad-domain.corp SRV TTL 0 100 389
ad-server.ad-domain.corp _ldap._tcp.Default-First-Site-Name._sites.dc._msdcs.ad-domain.corp SRV TTL
0 100 389 ad-server.ad-domain.corp
```

These records can be added automatically following the procedure below.

To prepare the AD server and AD domain for dynamic updates

1. Configure the master zone managing your AD domain

- a. Connect to SOLIDserver GUI.
- b. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- c. At the end of the line of the zone managing your AD domain, click on **⊞**. The properties page opens.
- d. In the panel **Access Control**, click on **[EDIT]**. The wizard opens: the page **Allow-query** appears.
- e. Click on **[NEXT]**. The page **Allow-transfer** appears.
- f. Click on **[NEXT]**. The page **Allow-update** appears.
- g. In the drop-down list **Type**, select *IP address*.
- h. In the drop-down list **Restriction**, select *Allow*.
- i. In the field **IP address**, type in the IP address of your AD server.
- j. Click on **[ADD]**, the IP address is visible in the list **ACL values**. Make sure that the IP address of the AD server was not denied in the list, denied hosts are preceded by (!).
- k. Click on **[OK]** to complete the operation. The report opens and closes. The IP address is visible in the panel *Access control*, in the list **Allow-update**.

2. Check the AD server configuration

- a. Connect to your AD server.
- b. Make sure the DNS resolver configured for the AD server is your EfficientIP DNS server. If not, you must set your EfficientIP DNS server as the resolver of your AD server.

3. Complete the AD server configuration

- a. Open a command prompt on the AD server.
- b. Send the AD server details - an A record - to your EfficientIP DNS server using the following command.

```
ipconfig /registerdns
```

- c. Add the missing key records to the DNS server using the command below.

```
nlttest /server:<ad-server-name>.<full-ad-domain> /dsregdns
```

The missing records link your AD server, AD domain and DNS servers.

4. Make sure the configuration is complete

- a. Go back to SOLIDserver GUI.
- b. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- c. Tick the zone managing your domain.
- d. In the menu, select **Edit > Status > Synchronize**. The wizard **Synchronization** opens.
- e. Click on **OK** to complete the operation. The report opens and closes, the page reloads.
- f. Click on the zone name, the page **All RRs** opens. The key records - 1 SOA, 2 A and 6 SRV - are listed and configured as expected.

When the AD server, DNS server and master zone managing your AD domain are configured properly, you can create a GSS-TSIG key and upload it to SOLIDserver.

Creating and Uploading the GSS-TSIG key

Once you prepared the AD server and AD domain for dynamic updates, you need to generate a GSS-TSIG key and upload it to your DNS server. The key must respect the following:

- The AD user named after the DNS server must generate the key.
- The key has the extension *.keytab* and is used by the DNS server to authenticate AD users.

You need to generate the key from the AD server and then upload it to SOLIDserver.

To generate the GSS-TSIG key on the AD server

1. Connect to your AD server using the credentials of an AD administrator.
2. Open a command prompt.
3. Go to the directory of your choice.
4. Generate the GSS-TSIG key following the command below, on one line.

```
ktpass -princ DNS/<name>.<domain>@<domain> -pass <password> -mapuser <name>@<domain> -ptype
KRB5_NT_PRINCIPAL +DesOnly -mapOp set -out <name>.<domain>.keytab

# <name> is the AD user sharing the same name as the DNS server.
# <domain> is the full AD domain: domain.TLD .
# <password> is the password of the AD user specified.
```

5. Copy the file in the directory of your choice for the upload.

Once you generated the key and saved it where you want, you must upload it to SOLIDserver. This upload is done on the page **All GSS-TSIG key**.

On the page, the keys are listed using their **Name**, **Encryption Type** and **Encryption number** in dedicated columns. You cannot edit the columns layout.

To upload the GSS-TSIG key to SOLIDserver

1. Connect to SOLIDserver GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
4. In the breadcrumb, click on **All GSS-TSIG keys**. The page refreshes.
5. In the menu, click on **+ Add**. The wizard **Add a GSS-TSIG Key** opens.
6. Click on **BROWSE** to find and upload your *.keytab* file.
7. Double-click on the name to select it.
8. In the field **File name**, the file is displayed.
9. Click on **OK** to complete the operation. The report opens and closes. The key is listed.

In the column **Name**, the uploaded key is displayed following the format declared during the generation: `DNS/<name>.<domain>@<domain>` .

You can delete a GSS-TSIG key if you no longer use it on any server or zone. For more details, refer to the section [Disabling Secure Dynamic Update](#).

Configuring your DNS server for Secure Dynamic Update

When the GSS-TSIG key is uploaded, you must configure the physical server managing your AD domain with it.

On the page *All servers*, the column **Multi-Status** indicates if the smart architecture has GSS-TSIG enabled but there is no GSS-TSIG key configured on the physical server(s) it manages.

If you manage your physical server via a smart architecture, you must first enable GSS-TSIG on the smart and then specify the *.keytab* file on the physical server.

To configure and specify the GSS-TSIG on a physical server managed via a smart

1. Enable GSS-TSIG on the smart server

- a. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
- b. At the end of the line of the smart server, click on . The properties page opens.
- c. Open the panel **GSS-TSIG** using  and click on **EDIT**. The wizard **Edit GSS-TSIG properties** opens.
- d. Tick the box **Use GSS-TSIG**.
- e. Click on **OK** to complete the operation. The report opens and closes. In the panel, the *GSS-TSIG status* is now **Enabled**. Now you must specify the GSS-TSIG key on the physical server.

2. Specify the GSS-TSIG key on the physical server

- a. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
- b. At the end of the line of the physical server managing your AD domain, click on . The properties page opens.

- c. Open the panel **GSS-TSIG** using  and click on **[EDIT]**. The wizard **Edit GSS-TSIG properties** opens
 - d. The box **Use GSS-TSIG** is ticked.
 - e. In the drop-down list **Select the key to use**, select the *.keytab* file you uploaded.
 - f. Click on **[OK]** to complete the operation. The report opens and closes. In the panel, the key details are displayed.
3. If you manage several physical servers via a Multi-Master architecture, you must repeat the actions of the step 2 for each EfficientIP DNS server.

If you manage your physical server outside a smart server, you must enable GSS-TSIG and specify the GSS-TSIG on the EfficientIP DNS server itself.

To configure a physical server with GSS-TSIG

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server managing your AD domain, click on . The properties page opens.
3. Open the panel **GSS-TSIG** using  and click on **[EDIT]**. The wizard **Edit GSS-TSIG properties** opens.
4. Tick the box **Use GSS-TSIG**. A drop-down list appears.
5. In the drop-down list **Select the key to use**, select the *.keytab* file you uploaded. If there is only one file on the page *All GSS-TSIG keys*, it is automatically selected.
6. Click on **[OK]** to complete the operation. The report opens and closes. In the panel, the key details are displayed.

Once you selected a *keytab* file on a physical server, the AD users can be authenticated, when they query the AD domain. Now you need to configure the AD users permissions and restrictions on the zone managing the AD domain.

Configuring your Zone for Secure Dynamic Update

When the physical server managing your AD domain is configured to authenticate AD users, you can configure the statement update-policy with permissions and restrictions that suit your needs directly on the zone managing your AD domain.

On the page *All zones*, the column **GSS-TSIG** indicates if the zones are configured with a GSS-TSIG key. The column has three possible values:

Table 42.1. Values of the column GSS-TSIG

Value	Description
N/A	The zone cannot be configured with GSS-TSIG.
No	The zone is not configured with GSS-TSIG: its update-policy statement has not been configured.
Yes	The zone is configured with GSS-TSIG: its update-policy statement is activated and configured with your entries and/or the default entries detailed below.

Before configuring the statement update-policy, keep in mind that:

- The first time you configure the statement, SOLIDserver automatically allows itself to edit the zone. The panel contains the policy you configured in the wizard and the following policy:

```
grant "ipmadmin" wildcard "*" ANY;
```

If your zone is managed via a view, SOLIDserver automatically adds the following policy:

```
grant "<view-name>" wildcard "*" ANY;
```

- If they suit your needs, we recommend that you configure the following update-policy entries:
 1. Allow the AD server to edit the master zone managing your domain. That way the AD server can advertise any hostname or IP address changes directly to the DNS server without configuring dynamic update again. The final policy should look as follows:

```
grant "<server-name>${<full-domain>" wildcard "*" ANY;
# In the GUI, this configuration is set as follows:
# Permission=grant, Identity=<server-name>${<full-domain>,
# Matchtype=wildcard, Tname=*, RR type=Any
```

2. Allow Windows computers registered in the AD domain to create their A and AAAA records in the zone.

```
grant "<full-ad-domain>" ms-self "<full-ad-domain>" specific A AAAA;
# In the GUI, this configuration is set as follows:
# Permission=grant, Identity=<server-name>${<full-domain>, Matchtype=ms-self,
# Tname=<full-domain>, RR type=Specific, Selected type(s)=A, AAAA
```

To configure secure dynamic update on the zone managing your AD domain

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. At the end of the line of the master zone managing your AD domain, click on **ⓘ**. The properties page opens.
3. In the panel **Access Control**, click on **[EDIT]**. The wizard opens: the page **Allow-query** appears.
4. Click on **[NEXT]**. The page **Allow-transfer** appears.
5. Click on **[NEXT]**. The page **Allow-update** appears.
6. Tick the box **Use GSS-TSIG/update-policy**. The page reloads and is now named **Update policy**.
7. Configure the dynamic update permissions and restrictions for your zone:

Table 42.2. Update-policy configuration parameters

Option	Description
Permission	Specify if you <i>grant</i> or <i>deny</i> update-policy rights. This field applies to all the other fields detailed below.
Identity	Specify the GSS-TSIG key sent by the client when they try to update the zone. It may look like <i><ad-server-name>\${<full-ad-domain></i> . Its use depends on the <i>Matchtype</i> you select, it can therefore partially depend on the <i>Tname</i> .
Matchtype	Select the matchtype that suits your needs: <i>subdomain</i> , <i>krb5-self</i> , <i>krb5-subdomain</i> , <i>ms-self</i> or <i>ms-subdomain</i> , described below. The matchtype only applies to the record type specified in the list Any . <p>subdomain allows to define the subdomain being updated. You must specify the domain in the field Tname following the format <i><domain>.<tld></i>, everything left of the specified <i><domain></i> becomes a match.</p> <p>wildcard allows to define the record being updated. You must specify the record's name in the field Tname using at least one wildcard (*). It can contain only a wildcard, in which case any record name can match.</p>

Option	Description
	krb5-self allows to define a rule based on Kerberos machine principal (host/QDN@REALM). The record being updated matches the QDN part of the Principal. The matching REALM must be specified exactly in the fields Identity and Tname .
	krb5-subdomain allows to define a rule based on Kerberos machine principal (host/QDN@REALM). The subdomain being updated matches the QDN part of the Principal. The matching REALM is what is specified the field Identity , or any subdomain of the specified Identity .
	ms-self allows to define a rule based on AD format principal (machinename\$@REALM) to update machinename.realm in the DNS. The matching REALM must be specified exactly in the fields Identity and Tname .
	ms-subdomain allows to define a rule based on AD format principal (machine-name\$@REALM) to update machinename.realm in the DNS. The matching REALM is what is specified the field Identity , or any subdomain of the specified Identity .
Tname	Specify a value to which applies the Matchtype <i>subdomain</i> , <i>krb5-self</i> , <i>krb5-subdomain</i> , <i>ms-self</i> or <i>ms-subdomain</i> . The expected format is detailed above, with each Matchtype.
RR type	Select which record type(s) the configuration set in the previous fields applies to. It can be <i>Specific</i> or <i>Any</i> .
<i>Specific</i>	Configure permissions for the record types of your choice via the lists Available types and Selected type(s) . In the list Available types , select <i>A</i> , <i>AAAA</i> , <i>CNAME</i> , <i>HINFO</i> , <i>AFSDB</i> , <i>MX</i> , <i>PTR</i> , <i>NS</i> , <i>SRV</i> , <i>TXT</i> , <i>WKS</i> , <i>NSAP</i> or <i>DNAME</i> and click on <input type="checkbox"/> to move it to the list Selected type(s) . Repeat this action for as many record types as you need.
<i>Any</i>	Allows to apply the configuration to all the update-policy record types: <i>A</i> , <i>AAAA</i> , <i>CNAME</i> , <i>HINFO</i> , <i>AFSDB</i> , <i>MX</i> , <i>PTR</i> , <i>NS</i> , <i>SRV</i> , <i>TXT</i> , <i>WKS</i> , <i>NSAP</i> and <i>DNAME</i> .

Once you configured the fields, click on . Your update-policy entry is moved to the **Update-policy list**. The page refreshes. Configure as many entries as you want.

To organize the list, use the buttons and . Each restriction or permission is reviewed and processed following the order set in the list.

- Click on to complete the operation. The report opens and closes. In the panel **Access control**, the list *Allow-update* is now replaced by the list **Update-policy**.

Even if you do not configure any entry in the **Update-policy list**, once you click on , SOLIDserver automatically adds the entry *grant "ipmadmin" wildcard "*" ANY;* that allows it to update your zone).

Each entry of the **Update-policy list** is listed on the page **User tracking**, in the column *Description* as follows: *DNS name: <name> Zone name: <name> Key: update-policy Value: <entry>* .

Once you configured the update policy statement, all the AD users querying your DNS server are authenticated by your AD server and only the allowed users can dynamically update the zone managing your AD domain.

Troubleshooting Secure Dynamic Update

If your configuration is not running as expected, you should make sure that:

- The AD server is running.
- The DNS server and AD server are at the same time and date.

You should configure an NTP server on both SOLIDserver and the AD server. For more details regarding NTP on SOLIDserver, refer to the section [Configuring NTP Servers](#).

3. The *.keytab* file you are using is the correct one.
 - a. Go the page All GSS-TSIG keys to make sure you uploaded the proper *.keytab* file. If not, upload it following the procedure in the section [Creating and Uploading the GSS-TSIG key](#).
 - b. Go to the properties page of the physical server managing the zone dedicated to your AD domain, in the panel GSS-TSIG the *.keytab* file is displayed. If the file is not the correct one, edit it following the procedure [To configure a physical server with GSS-TSIG](#).
4. The DNS server name is the same as the AD user declared in the *.keytab* file.

The AD user that generated it on the AD server and the physical DNS server managing your AD domain have the same name. If not, you might need to generate the *.keytab* file again following the section [Creating and Uploading the GSS-TSIG key](#). Then complete the physical server and zone configuration.

5. Your update-policy statement is properly configured. The order of the entries in the *Update policy list* is important: each restriction or permission is reviewed following the order you set in the list.

For more details, refer to the recommendations and procedure in the section [Configuring your Zone for Secure Dynamic Update](#).

Disabling Dynamic Update and Deleting Keys

If you no longer want users to dynamically update a zone you manage, you must:

- Edit the zone to empty the statement *allow-update* or *update-policy*.
- Make sure the server configuration suits your needs after your edited the zone.

Note that **you cannot delete TSIG or GSS-TSIG keys if they are used**. To delete them, you must first disable dynamic update or secure dynamic update.

Disabling Dynamic Update

To disable dynamic update, you must edit the zone to remove the relevant key from the statement *allow-update* and make sure that the ACL *admin* configuration of the DNS server suits your needs.

To disable dynamic update on a zone

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. Edit the zone:
 - a. At the end of the line of the zone of your choice, click on . The properties page opens.
 - b. In the panel **Access Control**, click on [\[EDIT\]](#). The wizard opens: the page **Allow-query** appears.
 - c. Click on [\[NEXT\]](#). The page **Allow-transfer** appears.
 - d. Click on [\[NEXT\]](#). The page **Allow-update** appears.
 - e. In the list **ACL values**, select an entry and click on [\[DELETE\]](#). It is no longer listed.
 - f. Repeat these actions for each entry in the list.

Keep in mind that if you delete all entries, the entry *admin* is automatically recreated at the top of the ACL list. It matches the configuration set on the server, which is why you must make sure that the ACL *admin* at server level suits your needs.

- g. Click on to complete the operation. The report opens and closes. The parameters are visible in the panel *Access control*, the list **Allow-update** is empty.
3. Check the server configuration:
- a. In the breadcrumb, click on the zone's server name. The server properties page opens.
 - b. In the panel **ACL**, in the list **DNS ACLs** select *admin*.
 - c. Click on . The wizard **ACL configuration Edit a DNS server** opens.
 - d. Edit the ACL configuration according to your needs using the fields *Type*, *Restriction* and *ACL values*.

Keep in mind that this ACL sets permissions and restrictions for all the zones of your server. So you must make sure that its content matches your needs.
 - e. Click on to complete the operation. The report opens and closes. The properties page is visible again.

Once you cleared the content of the statement *allow-update* and checked the ACL configuration of the server, you can remove TSIG keys following the section [Deleting TSIG and GSS-TSIG Keys](#).

Disabling Secure Dynamic Update

To disable secure dynamic update you must:

1. Edit the zone managing your AD domain. Go to the zone's properties page, edit the panel *Access control*, and on the page *Update policy* and untick the box **Use GSS-TSIG/update-policy**.
2. Edit the server managing your zone, it can be either the smart architecture managing your DNS physical server or the physical server itself if you manage it on its own.
 - Go the properties page of the relevant server, in the panel **GSS-TSIG** click on and disable GSS-TSIG: untick the box **Use GSS-TSIG** and click on .
 - From the same properties page, in the panel **ACL**: select *admin* and click on . Make sure the configuration suits your needs, if not make edit and click on to complete the operation.

Once you disabled dynamic update, you can delete the *.keytab* file following the procedure [To delete a GSS-TSIG key](#).

Deleting TSIG and GSS-TSIG Keys

You can delete DNS keys if:

- The TSIG key you want to delete from a server is no longer used by any of its zones. For more details, refer to the section [Disabling Dynamic Update](#).
- The GSS-TSIG key you want to delete from the page All GSS-TSIG keys is no longer used by any zone and GSS-TSIG is disabled on the server managing the zone. For more details, refer to the section [Disabling Secure Dynamic Update](#).

To delete a TSIG key

1. Make sure the key is no longer used in any statement of the server, view or zone.
2. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
3. At the end of the line of the server of your choice, click on . The properties page opens.
4. Open the panel **Keys** using .
5. In the list **DNS keys**, select the key you want to delete in the list **DNS keys**.
6. Click on . The wizard **Delete** opens.
7. Click on  to complete the operation. The report opens and closes. The properties page is visible again and the key is no longer listed.

To delete a GSS-TSIG key

1. Make sure you disabled secure dynamic update following the section [Disabling Secure Dynamic Update](#).
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
4. In the breadcrumb, click on **All GSS-TSIG keys**. The page refreshes.
5. Tick the key(s) you want to delete.
6. In the menu, click on  **Delete**. The wizard **Delete** opens.
7. Click on  to complete the operation. The report opens and closes. The properties page is visible again and the key is no longer listed.

Chapter 43. DNS Firewall (RPZ)

An *Response Policy Zone* (RPZ) is a set of resource records that associate domains, subdomains or IP addresses with specific response policies, based on domain data feeds provided by an external service or manually created by network administrators. This allows setting up a granular approach as, instead of blocking an entire domain, you can set exceptions or configure individual response rules for each subdomain.

In this sense, the RPZ is basically a DNS firewall that you can configure on the server: when a client queries a domain, subdomain or IP address listed in one of the RPZ zones managed by the server, the server replies with the related response policy. This mechanism is similar to an email anti-spam blacklist. In other words, it allows you to prevent DNS clients from accessing certain websites.

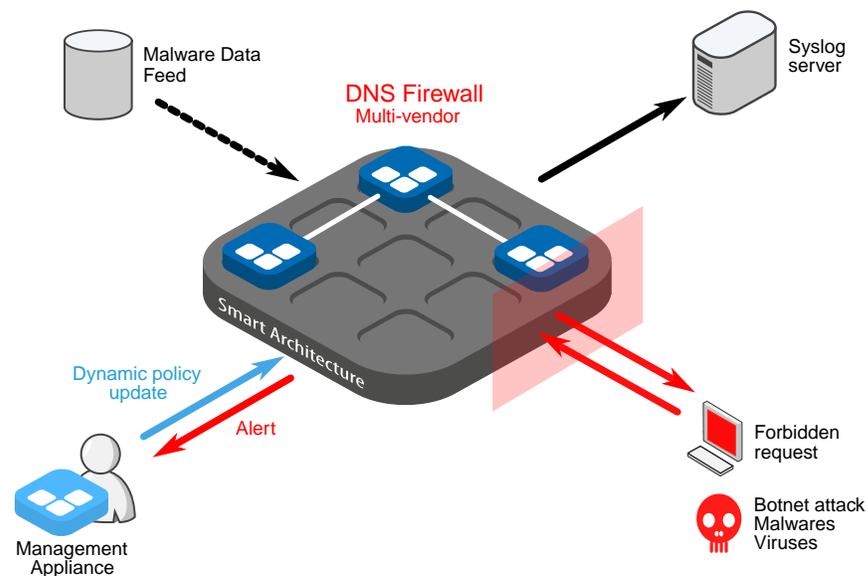


Figure 43.1. The DNS firewall

At zone level, you can decide which requests are redirected and where, as well as set a *NODATA* or even an *NXDOMAIN* response. The main benefit of this mechanism is that you can set up a filter using either a domain name or an IP address using the *CNAME*, *A* and *AAAA* records of the RPZ zone.

Within SOLIDserver, configuring the RPZ requires:

- 1. Managing an EfficientIP DNS server or a BIND DNS server.**

For more details, refer to the sections [Managing EfficientIP DNS Servers](#) and [Managing BIND DNS Servers](#).

- 2. Adding at least one RPZ Zone.**

For more details, refer to the section [Adding RPZ Zones](#) below.

- 3. Adding rules to the zone to implement specific policies depending on the queried data.**

For more details, refer to the section [Managing RPZ Rules](#) below.

Browsing RPZ Zones

Within the GUI, the RPZ zones are listed apart from regular DNS zones, on the page **All RPZ zones**. These zones contain records, used as rules, to define the relevant response policy depending on the queried data. They are all gathered and listed on the page **All RPZ rules**.

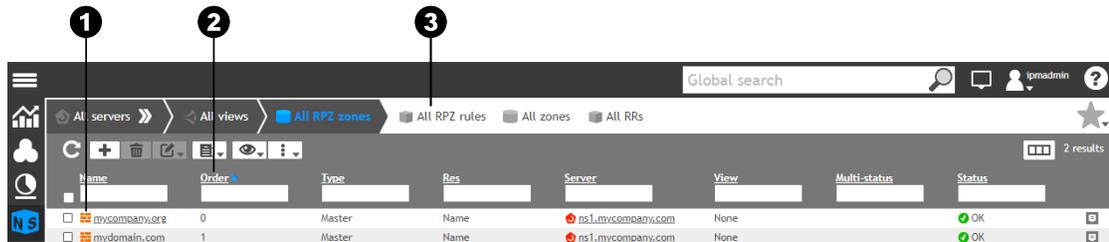


Figure 43.2. RPZ dedicated icons and page

The page *All RPZ zones* contains almost the same columns as the page *All zones*. However, note that:

- 1 Dedicated icon for the RPZ zones.
- 2 The column *Order* allows to display the position of the RPZ zone in the checking process, 0 being the first zone to be checked against the queried domain. As soon as a domain is found in one of the zone, the corresponding RPZ policy is applied and the following zones are ignored.
- 3 The page *All RPZ rules* lists all the records of RPZ zones. It even includes the SOA and NS records of each RPZ zone.

The RPZ configuration of a zone relies on RPZ rules. Once you added your BIND server to the page *All servers*, you need to add RPZ zones from the page *All RPZ zones* and configure RPZ rules through the addition of CNAME, A and AAAA records on the page *All RPZ rules*.

Browsing the RPZ Zones Database

The RPZ zones are managed on the page *All RPZ zones*. Their statuses are identical to regular zones. For more details, refer to the section [Understanding the DNS Zone Statuses](#).

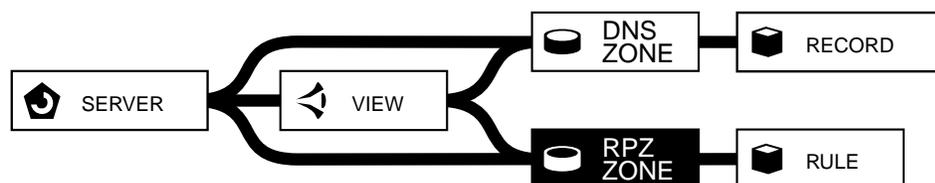


Figure 43.3. The RPZ zone in the DNS hierarchy

To display the list of RPZ DNS zones

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.

2. To display the RPZ zones of a specific server, in the column **Server** click on the name of the server of your choice. The page refreshes.

To display an RPZ DNS zone properties page

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. At the end of the line of the RPZ zone of your choice, click on **ⓘ**. The properties page opens.

Users of the group *admin* can create customized column layouts. The button **☰ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding RPZ Zones

Once you added your a BIND server to the DNS page All servers, you can add RPZ zones from the page All RPZ zones.

Keep in mind that:

- You can **only** configure RPZ **on Name zones**: it does not work on reverse zones.
- You can **only** configure RPZ **on Master zones or Slave zones**. Any other type of zone is irrelevant to the RPZ configuration.
- You **cannot use the name of an existing zone** when you create an RPZ zone, otherwise the domain name that you use can no longer be resolved using DNS.
- You **cannot add more than 32 RPZ zones** in one view or in one server if the server does not contain any views.

The zone name is returned in the DNS answer when using NODATA and NXDOMAIN policies. For security reasons, we recommend that you avoid using an obvious RPZ zone name, such as "RPZ-List", which could indicate that RPZ is implemented.

To add an RPZ zone

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the menu, click on **+ Add**. The wizard **Add a DNS zone** opens.
3. In the field **DNS server**, select a BIND server and click on **NEXT**. The next page of the wizard opens.
4. If you or your administrator created classes, the **DNS zone class** list is visible. Select a class or *None* and click on **NEXT**. The next page of the wizard opens.
5. In the list **DNS zone type**, select *Master* or *Slave*.
6. Click on **NEXT**. The next page of the wizard opens.
7. In the field **Name**, name your zone following the syntax given in RFC1034¹.
8. In the drop-down list **View**, select a view if you created any. If there are no views in the selected server, the list is empty.
9. The box **DNS firewall (RPZ)** is automatically ticked and displayed in gray.

¹Available on IETF website at <http://tools.ietf.org/html/rfc1034>.

10. In the drop-down list **Overriding rule**, *Given* is selected by default. For more details regarding the Overriding rules, refer to the section [Overriding RPZ Rules](#).
11. Click on **NEXT**. The last page of the wizard opens.
12. In the list **RPZ zones order**, the RPZ zones are sorted from the first one to be checked to the last one. To organize the list, select a zone and use the buttons  and  to move it.
13. If you are configuring a *Slave* zone:
 - a. Click on **NEXT**. The last page of the wizard opens.
 - b. Set up the list of master servers for the zone using the table below:

Table 43.1. DNS slave zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is required.
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

Once the IP, port and key are configured, click on **ADD**. The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to **DELETE** or **UPDATE** it once created.

14. Click on **OK** to complete the operation. The report opens and closes. The RPZ zone is listed, preceded by an orange rectangle and marked  *Delayed create* before being marked  **OK**.

Once you added a zone, a dedicated chart is available on the properties page of the physical server. For more details, refer to the section [Monitoring DNS Servers From their Properties Page](#).

You can configure their rules through records addition in the page All RPZ rules.

Editing RPZ Zones

Once you created an RPZ zone:

- You **cannot** rename it.
- You **cannot** reset its configuration and use it as a regular zone.

An RPZ zone can be edited on some level:

- You can edit its content and add as many RPZ rules as you please. For more details, refer to the section [Managing RPZ Rules](#).
- You can edit the class to the zone: add a new one, edit or remove it. For more details, refer to the chapter [Configuring Classes](#).
- You can configure an overriding rule or edit an existing one.
- You can edit the order in which the zones are checked.

To edit an RPZ zone

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. Right-click over the **Name** of the RPZ zone you want to edit. The contextual menu opens.
3. Click on **Edit**. The wizard **Edit a DNS zone** opens.
4. If you or your administrator configured classes at zone level, in the list **DNS zone class** you can select the class of your choice or *None*.
5. Click on **NEXT**. The next page of the wizard opens.
6. You cannot edit the **Name**, **View** (if the server contains any) and **DNS firewall (RPZ)** fields.
7. In the drop-down list **Overriding rule**, select the value of your choice: *Given*, *Disabled*, *Passthru*, *Nxdomain*, *Nodata* or *CNAME*. The page refreshes. For more details, refer to the section [Overriding RPZ Rules](#).
8. Click on **NEXT**. The next page of the wizard opens.
9. In the list **RPZ zones order**, the RPZ zones are sorted from the first one to be checked to the last one. To organize the list, select a zone and use the buttons **▲** and **▼** to move it.
10. If you are configuring a *Slave* zone:
 - a. Click on **NEXT**. The last page of the wizard opens.
 - b. Set up the list of master servers for the zone using the table below:

Table 43.2. DNS slave zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is required.
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

Once the IP, port and key are configured, click on **ADD**. The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to **DELETE** or **UPDATE** it once created.

11. Click on **OK** to complete the operation. The report opens and closes. The page *All RPZ zones* is visible again. The selected overriding rule is visible in the panel **Main properties** of the zone properties page.

Keep in mind that you can also edit an RPZ zone from its properties page:

- From the panel **Main properties**, you can edit the class and overriding rule of the zone.
- From the panel **Name servers**, you can edit the Authoritative DNS servers. For more details, refer to the section [Configuring Delegation at Zone Level](#).
- From the panel **Forwarding**, you can edit the zone forwarding parameters. For more details regarding the available parameters, refer to the section [Configuring DNS Forwarding at Server Level](#).
- From the panel **Notify**, you can edit IP addresses that should be notified of any changes on the master zone.

- From the panel **Access control**, you can set or edit allow-query, allow-transfer and allow-update options on your zones. For more details, refer to the section [Managing DNS Security](#).
- From the panel **Groups access**, the members of the admin group, can set and edit which groups should have or not the RPZ zone in their resources list.

Ordering RPZ Zones

When a user queries a server that manages several RPZ zones, SOLIDserver looks for a match in each RPZ zone, one after the other, following the configured order. The first matching rule is used and the other zones and rules are ignored.

You can order RPZ zones at any time. This allows, for instance, to set exceptions by placing zones configured by hand before those automatically updated by a 3rd party provider.

In addition, rules within a same zone respect some precedence order. For more details, refer to the section [Understanding the RPZ Rules Order](#).

To order RPZ zones

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. Right-click over the **Name** of any RPZ zone. The contextual menu opens.
3. Click on . The wizard **Edit a DNS zone** opens.
4. Click on **NEXT** until you get to the page **RPZ zones order**.
5. In the list **RPZ zones order**, the RPZ zones are sorted from the first one to be checked to the last one. To organize the list, select a zone and use the buttons  and  to move it.
6. Click on **NEXT** until you get to the last page of the wizard.
7. Click on **OK** to complete the operation. The report opens and closes. The page *All RPZ zones* is visible again. The RPZ zones position is visible in the column **Order**.

Converting RPZ Zones

Like any master or slave zone, RPZ zones can be converted from slave to master and vice versa.

To convert RPZ slave zones to master

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. Tick the slave zone(s) you want to convert.
3. In the menu, select  **Edit > Convert > To master**. The wizard **Zone conversion from slave to master** opens.
4. In the drop-down list **Remove old NS** you can decide what to do with the zone's NS records:
 - a. Select *Yes* to delete the NS record of the former slave zone.
 - b. Select *No* to keep them in the master zone for future use.
5. Click on **OK** to complete the operation. The report opens and closes. The page reloads. The converted zone is displayed on the page.

To convert RPZ master zones to slave

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. Tick the master zone(s) you want to convert.
3. In the menu, select **Edit > Convert > To slave**. The wizard **Zone conversion from master to slave** opens.
4. In the field **IP of master server**, specify the IP address of the master server that now has authority over the zone(s).
5. Click on **OK** to complete the operation. The report opens and closes. The page reloads. The converted zone is displayed on the page.

Deleting RPZ Zones

Deleting an RPZ zone can be done from the page All RPZ zones.

To delete an RPZ zone

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. Filter the list if need be.
3. Tick the RPZ zone(s) you want to delete.
4. In the menu, click on **Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The zone is marked **Delayed delete** until it is no longer listed.

Keep in mind that if you delete a zone that was configured with a Passthru overriding rule, any rule within the zone is deleted. Therefore, any record in the other zones of the server that matched one of the rules of the deleted zone is no longer configured with a Passthru exception. For more details, refer to the section [Overriding RPZ Rules](#).

Managing RPZ Rules

Within an RPZ zone the RPZ rules set policies for a specific query through the addition of *RPZ: <policy>* records. In fact, they create CNAME, A and AAAA records which syntax configures the filters of your choice.

SOLIDserver provides the configuration of four different policies that you can configure using requested domain names (QNAME) or IP addresses:

- **Redirection** is set with a record *RPZ: REDIRECT* on the page All RPZ rules. It allows to define which domain or IP address should be redirected toward which domain or IP address:

Domain name > domain name redirection

This redirection creates a CNAME record which name and value depend on the domain names stated during the configuration.

Domain name > IP address redirection

This redirection creates an A record if you redirect the domain name toward an IPv4 address or a AAAA record if you redirect the domain name toward an IPv6 address. This IP address can be the IP address of any equipment or even an entire subnet start address. Its name and value depend on the domain name and IP address stated during configuration.

IP address > domain name redirection

This redirection creates a CNAME record which name and value depend on the IP address and domain name stated during configuration.

IP address > IP address redirection

This redirection creates a CNAME record which name and value depend on the IP address stated during configuration.

- **NODATA** is set with a record *RPZ: NODATA*. It allows to set a NODATA response to any requested domain name or IP address. It basically creates CNAME record named after the domain name or IP address triggering this response.

Keep in mind that you can also set an NODATA policy using a Name Server Domain Name (NSDNAME) or Name Server IP address (NSIP).

- **NXDOMAIN** is set with a record *RPZ: NXDOMAIN*. It allows to set an denial of existence response to any requested domain name or IP address. It basically creates CNAME record named after the domain name or IP address triggering this response.

Keep in mind that you can also set an NXDOMAIN policy using a Name Server Domain Name (NSDNAME) or Name Server IP address (NSIP).

- **PASSTHRU** is set with a record *RPZ: PASSTHRU*. It allows to set an exception for the redirection or NODATA or NXDOMAIN response you set. It creates a CNAME record that, for instance, redirects **.domain.com* toward your company website but still grant access to the page *www.domain.com*.

Each RPZ rule record is created with a TTL of 3600 seconds. Once the policy applied, the TTL automatically drops to 5 seconds, following BIND behavior.

At server level, adding a rule to a zone adds the `response-policy` option in the `named.conf` file. SOLIDserver simply states in this option the RPZ zones managed by the server. In each of the RPZ zone file, the rules are listed as CNAME, A and AAAA records that respect the RPZ syntax.

Browsing the RPZ Rules Database

The RPZ zones contain records that set your rules. They are managed independently on the page *All RPZ rules*.

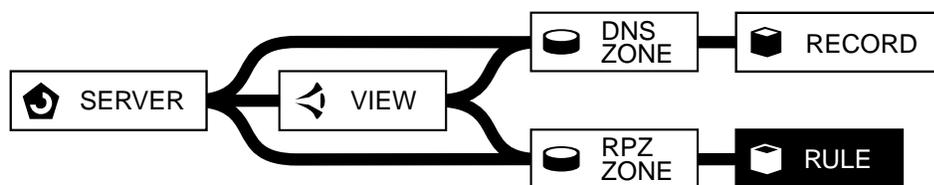


Figure 43.4. The RPZ rule in the DNS hierarchy

To display the list of RPZ rules

1. In the sidebar, go to **DNS > RPZ Rules**. The page **All RPZ rules** opens.
2. To display the list of rules of a specific RPZ zone, in the column **DNS zone name** click on the name of the zone of your choice. The page refreshes.

By default, the page *All RPZ rules* lists at least the SOA and NS record of each zone. Any additional record is listed and preceded by its own icon.

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the RPZ Rules Order

The RPZ is considered as the DNS *firewall* as RPZ rules are taken into account and implemented like firewall rules: once a match is found, it stops looking for other matches.

RPZ zones and RPZ rules respect a specific order:

1. The RPZ zones ordering matters

When a user queries a server that manages several RPZ zones, SOLIDserver looks for a match in each RPZ zone, one after the other, following the configured order. The first matching rule is used and the other zones and rules are ignored. For more details, refer to the section [Ordering RPZ Zones](#).

2. Within a single RPZ zone, rules respect a specific precedence

QNAME rules (i.e. domain name based rules) are preferred over IP based rules; IP rules are preferred over NSDNAME rules; NSDNAME rules are preferred over NSIP rules.

3. Within a single RPZ zone, name based rules follow a specific order

Among applicable QNAME or NSDNAME rules, the rule with the smallest name is preferred.

4. Within a single RPZ zone, IP based rules follow a specific order

- a. Among applicable IP or NSIP rules, the rule with the longest prefix length is preferred.
- b. Among IP or NSIP rules with the same prefix, the smallest IP address is preferred.

Configuring Rules Using Domain Names

A domain name (QNAME) can be used to set up a redirect, nodata, nxdomain and/or passthru response-policy through the addition of CNAME, A and AAAA records via the *Add an RPZ Rule* wizard.

Configuring a Redirection Using a Domain Name

The RPZ redirection policy can be configured using domain names. There are as many domain redirections as there are *RPZ: REDIRECT* records configured. You can either use a full domain name or specify some parts as variable, to include all the subdomains of a particular domain for instance.

To configure a redirection using a domain name

1. In the sidebar, go to  **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on  **Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *Domain*.

- In the field **Domain**, type in the domain name. It can be a full domain name or a partial one. For more details, refer to the table below.

Table 43.3. Domain name possible syntax when configuring an RPZ rule

Value	Description
domain.extension	The DNS client requesting this domain is redirected toward a domain name (refer to step 8) or toward an IP address (refer to step 9).
*.domain.extension ^a	The DNS client requesting any matching subdomain is redirected toward a domain name (refer to step 8) or toward an IP address (refer to step 9).
<value>.domain.extension	The DNS client requesting this specific name is redirected toward a domain name (refer to step 8) or toward an IP address (refer to step 9).

^aThe * (asterisk) is called the *wildcard* when used in front of a domain name.

- In the drop-down list **Policy**, select *Redirection*. You can set a redirection toward a domain name (refer to step 8) or toward an IP address (refer to step 9).
- Set the redirection toward the domain name of your choice:
 - In the drop-down list **Redirection target**, select *Domain*.
 - In the field **Target domain**, type in the target domain name of the redirection.
- Set the redirection toward the domain name of your choice:
 - In the drop-down list **Redirection target**, select *IPv4* or *IPv6*.
 - In the field **Target address**, type in the target IP address of the redirection, respecting the selected protocol version syntax.
- Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: REDIRECT* named after the source domain name, its **Value** is the target domain name or IP address depending on your configuration.

Configuring a NODATA Response Using a Domain Name

You can configure a NODATA response policy for clients requesting certain domain names. There is a NODATA response for as many domains as there are *RPZ: NO DATA* records configured.

To configure a NODATA response policy using a domain name

- In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
- In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
- In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
- In the drop-down list **Source**, select *Domain*.
- In the field **Domain**, type in the domain name. It can be a full domain name or a partial one. For more details, refer to the table below.

Table 43.4. Domain name possible syntax when configuring an RPZ rule

Value	Description
domain.extension	The DNS client requesting this domain gets a nodata response.

Value	Description
*.domain.extension	The DNS client requesting any matching subdomain gets a nodata response.
<value>.domain.extension	The DNS client requesting this specific name gets a nodata response.

- In the drop-down list **Policy**, select *Nodata*.
- Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NODATA* named after the source domain name, its **Value** is * following the BIND RPZ syntax in the zone file.

Configuring an NXDOMAIN Response Using a Domain Name

You can configure an NXDOMAIN response policy for clients requesting certain domain names. There is an NXDOMAIN response for as many domains as there are *RPZ: NXDOMAIN* records configured.

To configure an NXDOMAIN response policy using a domain name

- In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
- In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
- In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
- In the drop-down list **Source**, select *Domain*.
- In the field **Domain**, type in the domain name. It can be a full domain name or a partial one. For more details, refer to the table below.

Table 43.5. Domain name possible syntax when configuring an RPZ rule

Value	Description
domain.extension	The DNS client requesting this domain gets an nxdomain response.
*.domain.extension	The DNS client requesting any matching subdomain gets an nxdomain response.
<value>.domain.extension	The DNS client requesting this specific name gets an nxdomain response.

- In the drop-down list **Policy**, select *Nxdomain*.
- Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NXDOMAIN* named after the source domain name, its **Value** is . following the BIND RPZ syntax in the zone file.

Configuring a PASSTHRU Exception Using a Domain Name

Once you configured redirection and specific request responses, you can always configure a PASSTHRU exception for a particular domain name, subdomain, etc. There are as many domain name exceptions as there are *RPZ: PASSTHRU* records configured.

To configure a PASSTHRU response policy using a domain name

- In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
- In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
- In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.

4. In the drop-down list **Source**, select *Domain*.
5. In the field **Domain**, type in the domain name. It can be a full domain name or a partial one. For more details, refer to the table below.

You cannot use wildcard * when configuring a *passthru* from a domain name.

Table 43.6. Domain name possible syntax when configuring an RPZ rule

Value	Description
domain.extension	The DNS client requesting this domain gets a regular response.
<value>.domain.extension	The DNS client requesting this specific name gets a regular response.

6. In the drop-down list **Policy**, select *Passthru*.
7. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: PASSTHRU* named after the source domain name, its **Value** is *rpz-passthru* following the BIND RPZ syntax in the zone file.

Configuring Rules Using IP Addresses

An IP address can be used to set up a redirect, nodata, nxdomain and/or passthru response-policy through the addition of CNAME records via the wizard *Add an RPZ Rule*.

The RPZ follows a specific syntax similar to the reverse mapping (in-addr.arpa) in the zone file:

IPv4 rules display

Once created, the RPZ rules from an IPv4 address display the source IP address in reverse: *<prefixlength.B4.B3.B2.B1>*. In the zone file, the source IP address follows the RPZ syntax: *<prefixlength.B4.B3.B2.B1.rpz-ip>*.

IPv6 rules display

Once created, the RPZ rules from an IPv6 address display the source IP address in reverse: *<prefixlength.W8.W7.W6.W5.W4.W3.W2.W1>*. In the zone file, the source IP address follows the RPZ syntax: *<prefixlength.W8.W7.W6.W5.W4.W3.W2.W1.rpz-ip>*.

On the page All RRs, the column **Partial RR name** of RPZ records might contain *.zz. in .* It corresponds to *::* and allows you not to type in full the omitted *0000*: groups of the address.

Configuring a Redirection Using an IP Address

The RPZ redirection policy can be configured using a specific IPv4 or IPv6 address or range of addresses. There are as many IP addresses redirections as there *RPZ: REDIRECT* records configured.

Keep in mind that even though you can redirect a single address or a range of IP addresses (a subnet address for instance), the redirection target can only be one IP address.

To configure a redirection using an IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *IPv4* or *IPv6*.

5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select a prefix in the list. Your prefix might correspond to one IP address or to a range of IP addresses.
7. In the drop-down list **Policy**, select *Redirection*. You can set a redirection toward a domain name (refer to step 9) or toward an IP address (refer to step 10).
8. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *Domain*.
 - b. In the field **Target domain**, type in the target domain name of the redirection.
9. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *IPv4* or *IPv6*.
 - b. In the field **Target address**, type in the target IP address of the redirection, respecting the selected protocol version syntax.
10. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: REDIRECT* named after the source IP address and prefix displayed in reverse, its **Value** is the target domain name or IP address depending on your configuration.

Configuring a NODATA Response Using an IP Address

The RPZ NODATA policy can be configured using a specific IPv4 or IPv6 address or range of addresses. There are as many IP addresses redirections as there *RPZ: NODATA* records configured.

To configure a NODATA response policy using an IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *IPv4* or *IPv6*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select a prefix in the list. Your prefix might correspond to one IP address or to a range of IP addresses.
7. In the drop-down list **Policy**, select *Nodata*.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NODATA* named after the source IP address and prefix displayed in reverse, its **Value** is * following the BIND RPZ syntax in the zone file.

Configuring an NXDOMAIN Response Using an IP Address

The RPZ NXDOMAIN policy can be configured using a specific IPv4 or IPv6 address or range of addresses. There are as many IP addresses redirections as there *RPZ: NODATA* records configured.

To configure an NXDOMAIN response policy using an IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *IPv4* or *IPv6*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select a prefix in the list. Your prefix might correspond to one IP address or to a range of IP addresses.
7. In the drop-down list **Policy**, select *Nxdomain*.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NXDOMAIN* named after the source IP address and prefix displayed in reverse, its **Value** is *.* following the BIND RPZ syntax in the zone file.

Configuring a PASSTHRU Exception Using an IP Address

Once you configured the redirection and responses policies of your choice, the RPZ allows you to configure PASSTHRU exceptions for the IPv4 and IPv6 addresses or ranges of addresses of your choice. There are as many IP addresses exceptions as there are *RPZ: PASSTHRU* records.

To configure a PASSTHRU response policy using an IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *IPv4* or *IPv6*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select a prefix in the list. Your prefix might correspond to one IP address or to a range of IP addresses.
7. In the drop-down list **Policy**, select *Passthru*.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: PASSTHRU* named after the source IP address and prefix displayed in reverse, its **Value** is *rpz-passthru* following the BIND RPZ syntax in the zone file.

Configuring Rules Using Name Servers

In addition to filters through domain names and IP addresses, the RPZ filtering provides specific rule syntax that allows to provide alternate responses to queries made to authoritative name servers. You can set rules based on Name Server IP Address (NSIP) or Name Server Domain Name (NSDNAME). These filters add a extra suffix to the RPZ syntax and look as follows in the zone file: *<source-value>.rpz-nsip* or *<source-value>.rpz-nsdname* .

These records allow you to configure a redirection, an NXDOMAIN, a NODATA or a PASSTHRU response-policy to any query made to any zone managed by a Name Server whether you identified it through its IP address (NSIP) or through its domain name (NSDNAME).

Keep in mind that any of the zone managed by that authoritative Name Server are returned a NODATA or NXDOMAIN response if queried **EXCEPT** if you set a *passthru* exception for a particular zone or IP address managed by said Name Server. Indeed, as the NSDNAME and NSIP based rules are looked at last, if you set up a *passthru* based on a domain name (QNAME) or IP address, the *passthru* match is found before the name server domain name or IP address NODATA or NXDOMAIN policy. For more details, refer to the section [Understanding the RPZ Rules Order](#).

Configuring Rules using a Name Server Domain Name

A name server domain name can be used to set a NODATA or NXDOMAIN response to any query made to the zones it manages. This server name is usually embedded in the NS value of a domain name, once you retrieved it you simply need to add specify it as a Source Domain in the Add an RPZ Rule wizard.

To configure a redirection using a name server domain name

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSDNAME (domain name)*.
5. In the field **Domain**, type in the name server domain name.
6. In the drop-down list **Policy**, select *Redirection*. You can set a redirection toward a domain name (refer to step 8) or toward an IP address (refer to step 9).
7. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *Domain*.
 - b. In the field **Target domain**, type in the target domain name of the redirection.
8. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *IPv4 or IPv6*.
 - b. In the field **Target address**, type in the target IP address of the redirection, respecting the selected protocol version syntax.
9. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: REDIRECT* named after the name server domain name followed by the suffix *rpz-nsdname*, its **Value** is the target domain name or IP address depending on your configuration.

To configure a NODATA response policy using a name server domain name

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSDNAME (domain name)*.

5. In the field **Domain**, type in the name server domain name.
6. In the drop-down list **Policy**, select *Nodata*.
7. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NODATA* named after the name server domain name followed by the suffix *rpz-nsdname*, its **Value** is *** following the BIND RPZ syntax in the zone file.

To configure an NXDOMAIN response policy using a name server domain name

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSDNAME (domain name)*.
5. In the field **Domain**, type in the name server domain name.
6. In the drop-down list **Policy**, select *Nxdomain*.
7. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NXDOMAIN* named after the name server domain name followed by the suffix *rpz-nsdname*, its **Value** is *.* following the BIND RPZ syntax in the zone file.

To configure a PASSTHRU response policy using a name server domain name

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSDNAME (domain name)*.
5. In the field **Domain**, type in the name server domain name.
6. In the drop-down list **Policy**, select *Passthru*.
7. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NXDOMAIN* named after the source domain name followed by the suffix *rpz-nsdname*, its **Value** is *rpz-passthru* following the BIND RPZ syntax in the zone file.

Configuring Rules using a Name Server IP Address

The IP address of a name server can also be used to set a NODATA or NXDOMAIN response to any query made to the zones it manages. This server name IP address is usually embedded in the A glue record of the domain name NS record, once you retrieved it you simply need to add specify it as the Source Address with the prefix /32 in IPv4 and /128 in IPv6 in the Add an RPZ Rule wizard.

To configure a redirection using a name server IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.

2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSIP (IPv4)* or *NSIP (IPv6)*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select */32* for IPv4 or */128* for IPv6, if it was not automatically selected.
7. In the drop-down list **Policy**, select *Redirection*. You can set a redirection toward a domain name (refer to step 9) or toward an IP address (refer to step 10).
8. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *Domain*.
 - b. In the field **Target domain**, type in the target domain name of the redirection.
9. Set the redirection toward the domain name of your choice:
 - a. In the drop-down list **Redirection target**, select *IPv4* or *IPv6*.
 - b. In the field **Target address**, type in the target IP address of the redirection, respecting the selected protocol version syntax.
10. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: REDIRECT* named after the source IP address and prefix in reverse followed by the suffix *rpz-ip*, its **Value** is the target domain name or IP address depending on your configuration.

To configure a NODATA response policy using a name server IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSIP (IPv4)* or *NSIP (IPv6)*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select */32* for IPv4 or */128* for IPv6, if it was not automatically selected.
7. In the drop-down list **Policy**, select *Nodata*.
8. Click on to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NODATA* named after the source IP address and prefix displayed in reverse followed by the suffix *rpz-ip*, its **Value** is *** following the BIND RPZ syntax in the zone file.

To configure an NXDOMAIN response policy using a name server IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.

3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSIP (IPv4)* or *NSIP (IPv6)*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select */32* for IPv4 or */128* for IPv6, if it was not automatically selected.
7. In the drop-down list **Policy**, select *Nxdomain*.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: NXDOMAIN* named after the source IP address and prefix displayed in reverse followed by the suffix *rpz-ip*, its **Value** is *.* following the BIND RPZ syntax in the zone file.

To configure a PASSTHRU response policy using a name server IP address

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *NSIP (IPv4)* or *NSIP (IPv6)*.
5. In the field **Address**, type in the IP address following the appropriate syntax.
6. In the drop-down list **Prefix**, select */32* for IPv4 or */128* for IPv6, if it was not automatically selected.
7. In the drop-down list **Policy**, select *Passthru*.
8. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: PASSTHRU* named after the source IP address and prefix displayed in reverse followed by the suffix *rpz-ip*, its **Value** is *rpz-passthru* following the BIND RPZ syntax in the zone file.

Configuring Other Rules

An administrator might need to configure rules that do not use a domain name, IP address, name server domain name or IP address as a source; in which case, you need to specify yourself the partial RR name, as the full name is automatically created as follows: *<partial-rr-name>.<zone-name>*. Keep in mind that the procedure below is advanced configuration and therefore, the consistency of the data in the field *Value* of the wizard is not checked. So if the syntax does not comply with RPZ, the filter it sets should obviously not work.

For these sources, the available rules are the same: redirection, Nodata, Nxdomain or Passthru.

To configure an RPZ rule using a specific source

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add an RPZ Rule** opens.
4. In the drop-down list **Source**, select *Other*.
5. In the field **Value**, type in the source identification following the RPZ syntax (with the appropriate values and suffixes).

6. In the drop-down list **Policy**, select the policy that suits your needs. If you select *Redirection*, you need to specify a domain name or an IP address.
7. Click on **OK** to complete the operation. The report opens and closes. The record is now listed. The column **RR name** displays an *RPZ: <policy>* named after the content of the **Value** field in the wizard, its value depending on your configuration.

Deleting Rules

At any time, you can delete a rule. In other words, you can delete an RPZ record. In the procedure below, we delete record within a specific RPZ zone but you can also delete records from the page *All RPZ rules* without filtering a specific zone.

To delete an RPZ rule

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the column **Name**, click on the name of the RPZ zone of your choice. The page **All RPZ rules** opens.
3. Filter the list if need be.
4. Tick the RPZ record(s) you want to delete.
5. In the menu, click on **Delete**. The wizard **Delete** opens.
6. Click on **OK** to complete the operation. The report opens and closes. The record is marked **Delayed delete** until it is no longer listed.

Overriding RPZ Rules

Once you created an RPZ zone, you can override all the rules it contains using the drop-down list *Overriding rule* in the RPZ zone addition and edition wizards.

Overriding Rules Specificities

- There are six overriding rules that are named after the existing policies: *Given*, *Disabled*, *Nxdomain*, *Nodata*, *CNAME* and *Passthru*.
- The overriding rules must be set a zone level.
- The overriding rules apply to all the rules of the zone except the *Passthru* that applies to the server: once set on an RPZ zone, the *Passthru* applies to all the rules of the zone as well as any matching RPZ rule name within the server.

To override the RPZ rules when creating an RPZ zone

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.
2. In the menu, click on **Add**. The wizard **Add a DNS zone** opens.
3. In the field **DNS server**, select a BIND server and click on **NEXT**. The next page of the wizard opens.
4. If you or your administrator created classes, the **DNS zone class** list is visible. Select a class or *None* and click on **NEXT**. The next page of the wizard opens.
5. In the list **DNS zone type**, select *Master* or *Slave*.
6. Click on **NEXT**. The last page of the wizard opens.

7. In the field **Name**, name your zone following the syntax given in RFC1034 ².
8. In the drop-down list **View**, select a view if you created any. If there are no views in the selected server, the list is empty.
9. The box **DNS firewall (RPZ)** is automatically ticked and displayed in gray.
10. In the drop-down list **Overriding rule**, select the value of your choice according to the table below.

Table 43.7. Overriding rules

Field	Description
Given	All the rules specified in the zone are applied. It is the default value.
Disabled	Disables all the RPZ rules of the zone.
Nxdomain	Applies an NXDOMAIN response to all the RPZ rules of the zone, no matter the policy configured when creating them.
Nodata	Applies an NODATA response to all the RPZ rules of the zone, no matter the policy configured when creating them.
CNAME	Applies a redirection to all the RPZ rules of the zone. Once selected, the field <i>Domain</i> appears: type in the FQDN of your choice. All the RPZ rules of the zone are ignored and redirected toward the specified domain.
Passthru	Applies a PASSTHRU exception to all the RPZ rules of the zone. Keep in mind that if any RPZ rule on the server matches the name of one or more of the RPZ rules of your zone, setting the passthru policy on your zone applies to these rules as well. The passthru policy applies to the RPZ rules of a server and not only of a specific zone.

The page refreshes.

11. If you are configuring a *Slave* zone:
 - a. Click on **NEXT**. The last page of the wizard opens.
 - b. Set up the list of master servers for the zone using the table below:

Table 43.8. DNS slave zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is required.
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

Once the IP, port and key are configured, click on **ADD**. The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to **DELETE** or **UPDATE** it once created.

12. Click on **OK** to complete the operation. The report opens and closes. The page *All RPZ zones* is visible again. The selected overriding rule is visible in the panel **Main properties** of the zone properties page.

To override the RPZ rules when editing an RPZ zone

1. In the sidebar, go to **DNS > RPZ Zones**. The page **All RPZ zones** opens.

²For more details, refer to <http://tools.ietf.org/html/rfc1034>, on page 7.

2. Right-click over the **Name** of the RPZ zone you want to edit. The contextual menu opens.
3. Click on . The wizard **Edit a DNS zone** opens.
4. The field **Name** displays the name of the zone. It is in gray: you cannot edit it.
5. The box **DNS firewall (RPZ)** is ticked and displayed in gray.
6. In the drop-down list **Overriding rule**, select the value of your choice. For more details, refer to the table [Overriding rules](#). The page refreshes.
7. If you are configuring a *Slave* zone:
 - a. Click on . The last page of the wizard opens.
 - b. Set up the list of master servers for the zone using the table below:

Table 43.9. DNS slave zone parameters

Field	Description
Master IP address	Type in the master server IP address. This field is required.
Port	You can type in the number of the port dedicated to communicating with the slave zone. This field is optional.
TSIG key	In the drop-down list, select a TSIG key set at server level. This field is optional. For more details, refer to the section Securing the Management of DNS Servers Using a TSIG Key .

Once the IP, port and key are configured, click on . The configuration is listed in the list **Masters**. Repeat these actions for as many servers as needed. You can select a master in the list to  or  it once created.

8. Click on  to complete the operation. The report opens and closes. The page *All RPZ zones* is visible again. The selected overriding rule is visible in the panel **Main properties** of the zone properties page.

RPZ Limitations

- You **can implement the RPZ only on BIND servers**. Within the GUI this implies remotely managing a BIND server or creating an EfficientIP DNS server.

For more details, refer to the section [Managing EfficientIP DNS Servers](#) or [Managing BIND DNS Servers](#).

- You **cannot add more than 32 RPZ zones**:
 - In one view of a BIND server.
 - On a BIND server that does not contain any views.
- You **cannot set the order of the rules**. Therefore the order of the zones is important.

The RPZ rules are applied one after the other respecting both the rule type and their order within each zone, and views, of the server. Once a queried record is matched in one zone, it is ignored in the following rules of the zone. For more details regarding the order of the rules within a zone, refer to the section [Understanding the RPZ Rules Order](#).

Chapter 44. Hybrid DNS Service

SOLIDserver provides a Hybrid DNS service that reduces corruption risks for BIND DNS engines. Hybrid DNS incorporates an alternative DNS engine based on NLnet Labs Unbound and NSD engines that can automatically switch from standard BIND service to Hybrid if their configuration is compatible.

Depending on your configuration, **authoritative engines switch to BIND/NSD hybrid** and **recursive engines switch to BIND/Unbound hybrid**. Note that you cannot decide to switch to NSD or Unbound, the switch is automatic and entirely dependent on the engine configuration.

Once the switch is complete, the DNS engine footprint is more complex to analyze and less prone to malicious attacks as the DNS mechanism is different, it avoids BIND security flaws altogether. Therefore, in the event of an attack or important security issue, the switch to Hybrid ensures data security and avoids its potential corruption.

Keep in mind that Hybrid engines have some [limitations](#) compared to BIND engines.

Checking the Compatibility with Hybrid

Checking the compatibility with Hybrid implies to:

1. Match the basic Hybrid requirements.
2. Check that no parameter set at server or zone level is incompatible with Hybrid.
3. Generate the incompatibility report, if need be.
4. Edit the server configuration to make sure that none of the parameters set are incompatible with Hybrid.

Before switching, you need to understand that you cannot decide if your physical server switches to BIND/NSD or BIND/Unbound. As a general rule, if your server is compatible with Hybrid, the following switch occurs:

- If the smart server **recursion** is set to **yes**, a Hybrid compliant server can **switch to BIND/Unbound**.
- If the smart server **recursion** is set to **no**, a Hybrid compliant server can **switch to BIND/NSD**.

Matching Hybrid Basic Requirements

The first step toward switching to Hybrid is to match the following **Hybrid basic requirements**:

- You can only convert servers to Hybrid from SOLIDserver hardware or software appliance.
- The servers you want to switch must be EfficientIP DNS servers.
- The servers you want to switch must be managed via a smart architecture. The changes are pushed to the physical server.
- The smart architecture cannot be compatible with Hybrid if it does not manage only BIND servers.
- The physical server status must be  OK, you cannot switch a server in *Timeout*.

On the DNS All servers list, the **Hybrid DNS compatibility** and **Forced Hybrid DNS compatibility** columns allow you to see if you can switch your BIND physical servers.

In addition, the **Multi-status** column at server, view, zone and RR level provides you with all the potential incompatibilities with Hybrid. For more details, refer to the section [Understanding the Column Multi-Status](#). For more details regarding how to change a page listing template, refer to the section [Customizing the List Layout](#).

This information is also provided on the smart architecture edition wizard: the *Compatible with a Hybrid DNS Engine* field indicates the Hybrid compatibility of the physical servers managed.

Making Sure the Server Configuration is Compatible with Hybrid

If the smart architecture managing your physical server, is marked *No* in the Hybrid DNS compatibility column, the physical server cannot be switched to Hybrid. **If the server is set with one of the following options and configurations, it cannot be switched to Hybrid:**

- The DNS server type is different from a SOLIDserver Hardware or Virtual Appliance EfficientIP DNS server (for instance a server using packages, an agentless server, a generic server, etc.).
- The server contains views.
- The server contains zones other than master, slave, forward or stub.
- The server contains master and/or slave zones as well as forward and/or stub zones. With Hybrid, the server is either only authoritative or only recursive.
- One or more server zones are RPZ compliant.
- One or more server zones are signed with DNSSEC.
- The server configuration combines authoritative and recursive zones:
 - If the DNS recursion set to yes and the server contains master and/or slave zones, the server cannot switch to Hybrid.
 - If the DNS recursion set to no and the server contains forward and stub zones, the server cannot switch to Hybrid.
 - If the DNS recursion set to yes with TSIG keys.

You must change your configuration to match Hybrid requirements if you want to switch to Hybrid. During the switch, SOLIDserver checks once more all the parameters to make sure that your server is compatible once more.

If you want to have a complete list of all the parameters and options that need to be edited, refer to the section [Generating the Hybrid Incompatibilities Report](#) below.

Generating the Hybrid Incompatibilities Report

If the smart architecture managing your BIND servers is not compatible with Hybrid, you can generate the List Hybrid DNS Engine incompatibilities report to have a detailed list of all the parameters that do not comply with hybrid following the procedure below.

To generate the Hybrid DNS Engine incompatibilities report

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. Tick the smart server managing the physical server you intend to switch to Hybrid.
3. In the menu, select  **Report > Hybrid DNS Engine incompatibilities**. The wizard **Hybrid incompatibilities report** opens.
4. In the Report format list, select HTML or PDF.

5. Click on **NEXT**. The last page of the report opens.
6. In the drop-down list **Action**, select the kind of report to want to generate.

Table 44.1. Hybrid report available actions

Field	Description
Generate new data	This action generates a report that lists all to the incompatibilities with Hybrid at the moment you create it. Once generated, this report is available in the list named as follows: <date time>.
Schedule the report	This action generates a graph as regularly as you need.

7. If you chose to **Schedule the report**, configure the reports using these fields.

Table 44.2. Scheduled report parameters

Field	Description
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, <i>Every hour</i> is selected.
Minute	Select a period of time, minutes-wise. By default, <i>Every minute</i> is selected.
Name	A default name is already filled in, you can edit this scheduled export name if you want.
Mail to	In this drop-down list, select the group which users should receive the export notification email. This email cannot be sent if the users email address is not valid or if your SMTP relay is not configured. For more details, refer to the section Configuring the SMTP Relay . By default, the first of your groups, in the ASCII alphabetic order, is selected.
Rights as	Select a user. His/her rights and limitations are applied in the report: only the items this user has access to are listed in the export.

8. Click on **OK** to generate the report. The page *Report* opens and works for a while.
9. You can click on **DOWNLOAD** to save the report immediately.

When the report is generated, it is available on the page *Reports*. For more details, refer to the procedure [To list the reports](#).

Once you generated the report, all the parameters that are not compatible with Hybrid are listed and you need to correct them all until your smart server is marked compatible. You can generate as many reports as you want, every report is available on the Reports page of the Administration module.

To find the Hybrid DNS Engine incompatibilities report

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Reports**. The page **Reports** opens.
3. In the list, the Hybrid DNS incompatibilities report are listed.

Once the physical server is Hybrid compliant, on the page All servers, the column **Hybrid DNS compatibility** is marked *Yes* and the smart architecture edition wizard field *Compatible with a Hybrid DNS Engine* is also marked *Yes*.

Switching to Hybrid DNS

Once your smart architecture is compatible with Hybrid, you can switch it. If your server is not compatible with Hybrid, you need to change its configuration as some parameters might prevent the switch. For more details, refer to the section [Checking the Compatibility with Hybrid](#).

The architecture can contain one or several BIND servers that you can all switch. Keep in mind that, if you only switch one server, the other servers share the same limitations that the Hybrid servers. So, before switching to Hybrid you should probably make sure that none of its limitations prevent you from using your server with all the parameters you usually need. For more details, refer to the section [Hybrid DNS Engines Limitations](#).

The switch to Hybrid actually follows this order:

1. All the Hybrid incompatibilities checks are made again.
2. If the server is actually compatible, the relevant Hybrid configuration is pushed to the physical server.
3. Once the whole configuration is successfully pushed, BIND service is disabled and stopped and the relevant Hybrid service (NSD or Unbound) is enabled and started.

In some rare cases, you might have a Hybrid server listed among your servers outside a smart architecture. As you cannot manage a Hybrid server outside a smart architecture, you need to switch it to BIND, add it to your smart architecture and then switch it again to Hybrid. For more details, refer to the procedures [To switch a physical server from Hybrid to BIND DNS](#) and [To switch a physical server from BIND to Hybrid DNS](#).

To switch a physical server from BIND to Hybrid DNS

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Make sure the BIND physical server you want to switch to Hybrid belongs to a smart architecture compatible with Hybrid.
3. Make sure the server status is **OK**.
4. Tick the physical server you want to switch.
5. In the menu, select **Tools > Expert > Switch DNS Engine > To NSD / Unbound**. The **Switching the DNS Engine** wizard opens.
6. Click on **OK** to complete the operation. The report opens and works until the relevant DNS service restarts. The physical server **Status** is **OK** and its **Version** indicates the engine name it switched to.

Your server configuration switches to Unbound or NSD on its own, based on its configuration. Once the switch is complete, the compatibility with Hybrid is forced: this implies that a set of configurations can no longer be set. For more details regarding NSD or Unbound specificities and limitations, refer to the sections [The Server Switched to NSD](#) and [The Server Switched to Unbound](#) below. As for the Hybrid limitations in general, refer to the section [Hybrid DNS Engines Limitations](#).

To display the Hybrid engine the server switched to in the DNS module

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the column **Version**, the engine and version are displayed.

Like any other server, you can check on a Hybrid server from the columns *Status* and *Sync*. For instance, make sure that the smart architecture can push your configuration on the physical server, if not the smart is marked *Locked synchronization*. For more details regarding this status, refer to the section [Handling the Status Locked Synchronization](#).

To display the Hybrid engine the server switched to in the Administration module

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. Next to the **DNS server**, the engine that runs is indicated between brackets.

From the Services configuration page, you can enable, disable, stop and start the Hybrid DNS server. For more details, refer to the section [Handling Services](#).

In the same way, from this page you can download the NSD or Unbound configuration file depending on which one is running. For more details refer to the section [Downloading the DNS/DHCP/DHCPv6 Configuration File](#).

The Server Switched to NSD

NSD Engines are designed to manage authoritative DNS configurations. Once the switch was successfully performed, a set of BIND options and configurations is emulated to suit NSD requirements.

However, you should be aware of a set of NSD engines specificities and limitations that shape the configuration that you can or cannot set from the GUI.

NSD engines specificities

- NSD servers are exclusively authoritative: only master and slave zones are supported.
- All records handled by BIND are handled by NSD but DNSSEC records are only supported if you use the BIND/NSD server as slave server in a Stealth or Master/Slave smart architecture.
- Each change made to the server or zones creates a new NSD configuration or zone file, copies the former files and pushes the new configurations on the physical server.
- Every change made to the records database rebuilds the NSD database and creates a new zone to ensure that the changes are pushed to the physical server as soon as possible.

NSD engines limitations

- You cannot create *forward*, *stub*, *hint* or *delegation-only* zones on an NSD server.
- Not all ACLs are supported:
 - *none*, *any*, *localhost* and all the access control lists based on IP or network addresses are supported.
 - The *localnets* ACL is ignored.
- The allow-transfer and allow-notify clauses set on your BIND server are converted as follows after a switch to NSD:
 - If the allow-transfer clause is not specified at server or zone level, a default configuration is pushed on the NSD server to allow any user to transfer master and slave zones via AXFR.

- If the allow-notify clause is not specified at server or zone level, the clause value on the NSD server is set to respect BIND default behavior and allow proper synchronization of the master and slave zones.
- NSD supports all the RRL settings except the Log only option. For more details, refer to the section [Limiting the Number of Responses of a Server](#).

The Server Switched to Unbound

Unbound Engines are designed to manage recursive DNS configurations. Once the switch was successfully performed, a set of BIND options and configurations are emulated to suit Unbound requirements. However, you should be aware of a set of Unbound engines specificities and limitations that shape the configurations that you can or cannot set from the GUI.

Unbound engines specificities

- Unbound servers are exclusively recursive: only forward and stub zones are supported.
- BIND statements are interpreted as follows:
 - If the allow-recursion is specified on BIND, its value is used to set the allow-query statement on Unbound.
 - If the allow-recursion is not specified on BIND, the localhost is set on Unbound.
- ACLs are only supported to configure the allow-recursion statement only at server level. For more details regarding ACLs, refer to the section [Unbound Engines Limitations](#) below.
- On forward zones, the forward parameter can only be set to *first*.
- If the BIND server is configured with the forward parameter (set to any value but *none*) and forwarders, the switch to Hybrid DNS creates a forward zone named "." that emulates all specified parameters. Keep in mind that if a "." forward zone already exists, the list of forwarders of both zones are merged into one. Other parameters of the existing "." forward zone are ignored.

Unbound Engines Limitations

- You cannot create *master*, *slaver*, *hint* or *delegation-only* zones on an Unbound server. Converting master zones to slave on Unbound servers deletes all the records of the original master zone.
- Not all ACLs are supported:
 - *none*, *any*, *localhost* and all the access control lists based on IP or network addresses are supported.
 - The *localnets* ACL and TSIG keys are not supported.
- Stub zones cannot be configured with:
 - *forward* and parameter *forwarders*.
 - *stub-first* and parameter *stub-primes*: they do not have any equivalent in BIND.
- Forward zones cannot be configured with the forward parameter set to *only*.
- Unbound handles the edns-udp-size option in a unique way:
 - If the option was set before switching, the specified value is set on the Unbound *ipv4-edns-size* and *ipv6-edns-size* options. Keep in mind that in this case, *ipv4-edns-size* has precedence over *ipv6-edns-size*.
- Unbound does not support RRL. For more details regarding RRL, refer to the section [Limiting the Number of Responses of a Server](#).

Hybrid DNS Engines Limitations

Once you switched your DNS service to Hybrid, you can configure and manage it through a smart architecture. However, Hybrid has some limitations:

- It is impossible to import a Hybrid configuration.
- No statistics regarding Hybrid servers are retrieved, therefore the server properties page does not contain any graph.
- You can only switch to Hybrid physical servers managed via a smart architecture
- After a fresh installation, the service default type is BIND. You need to manage the server through a smart architecture and then switch it.
- Only the options compatible with BIND are supported: any hybrid vendor option that does not have any counterpart in BIND cannot be set through SOLIDserver.
- The RNDC commands are not supported: you cannot perform the commands force-notify, force-refresh, force-retransfer, querylog and flush cache on Hybrid compliant servers.
- The options inheritance is not supported per se. However, after switching to Hybrid, your server configuration is directly set at zone level.
- SOLIDserver does not retrieve data from a Hybrid server. However, if you manage a Hybrid server via a smart you can synchronize the architecture to push any changes made from the GUI to the server (content or configuration file) from the smart to the physical server.
- Any change made to a Hybrid server restarts the service.
- Dynamic Update is not supported by Hybrid.
- ACL use is limited:
 - All ACLs based on IP and network addresses are supported.
 - The *any*, *localhost* and *none* ACL are supported in their IP address form.
 - The *localnet* ACL is not supported.
- Views are not supported.
- RPZ zones are not supported.
- Signing zones with DNSSEC is only possible if the BIND/NSD Hybrid server is managed via a smart architecture in which the BIND/NSD server is a slave which master is a BIND server. Indeed, NSD servers cannot be signed in DNSSEC but as slave servers they can handle DNSSEC records.

Forcing Compatibility with Hybrid

To provision a switch to Hybrid, you can force the compatibility with Hybrid on smart architectures. This action allows you to make sure that all the parameters and configurations you set on your server (at server and/or zone level) respect Hybrid requirements for BIND/NSD or BIND/Unbound. That way, you can switch your engine right away and do not need to perform any configuration changes, whether you were planning to on a particular day or because a CVE release impacts your BIND servers security.

To force the compatibility with Hybrid

1. In the sidebar, go to  **DNS** > **Servers**. The page **All servers** opens.

2. Make sure the BIND physical server you want to switch to Hybrid belongs to a smart architecture compatible with Hybrid.
3. Right-click over the **Name** of the smart architecture that manages this server. The contextual menu opens.
4. Click on . The **Edit a DNS server** wizard opens.
5. If you are editing a *Master/Slave*, *Stealth*, *Multi-Master* or *Single-Server* architecture follow the steps below:
 - a. Click on **NEXT** until you get to the page **DNS servers role configuration**.
 - b. Tick the box **Expert mode**.
 - c. Click on **NEXT**. The page **Advanced settings** opens.
 - d. Tick the box **Force Hybrid DNS compatibility**.
6. If you are editing a *Farm* architecture follow the steps below:
 - a. Click on **NEXT** until you get to the page **Advanced settings**.
 - b. Tick the box **Force Hybrid DNS compatibility**.
7. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is marked *Yes* in the *Forced Hybrid DNS compatibility*.

Switching Back to BIND

As Hybrid engines imply a set of limitations that might prevent you from configuring your DNS server according to your needs, mixing authoritative and recursive zone for instance, you can switch back to BIND. As all the NSD and Unbound options that you can set from the GUI have an equivalent in BIND, switching the engine back to BIND can be performed at any time.

To switch a physical server from Hybrid to BIND DNS

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. Make sure the BIND physical server you want to switch to Hybrid belongs to a smart architecture compatible with Hybrid.
3. Make sure the server status is  **OK**.
4. Tick the physical server you want to switch.
5. In the menu, select . **Tools > Expert > Switch DNS Engine > To BIND**. The **Switching the DNS Engine** wizard opens.
6. Click on **OK** to complete the operation. The report opens and works until the relevant DNS service restarts. The physical server **Status** is  **OK** and its **Version** indicates it switched to BIND.

Once you switched a Hybrid server engine to BIND, the option *Force Hybrid DNS compatibility* is still enabled. To be able to configure the BIND server without the Hybrid limitations, you must edit the smart architecture and untick the box.

To remove the forced compatibility with Hybrid

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.

2. Make sure that the smart architecture you want edit does not manage any Hybrid server.
3. Right-click over the **Name** of the smart architecture. The contextual menu opens.
4. Click on . The **Edit a DNS server** wizard opens.
5. If you are editing a *Master/Slave*, *Stealth*, *Multi-Master* or *Single-Server* architecture follow the steps below:
 - a. Click on **NEXT** until you get to the page **DNS servers role configuration**.
 - b. Tick the box **Expert mode**.
 - c. Click on **NEXT**. The page **Advanced settings** opens.
 - d. Untick the box **Force Hybrid DNS compatibility**.
6. If you are editing a *Farm* architecture follow the steps below:
 - a. Click on **NEXT** until you get to the page **Advanced settings**.
 - b. Untick the box **Force Hybrid DNS compatibility**.
7. Click on **OK** to complete the operation. The report opens and closes. The smart architecture is marked *No* in the *Forced Hybrid DNS compatibility*.

Administrating the Backup and Restoration of Hybrid Configurations

As Hybrid DNS engines differ from BIND *named* engine, there are a set of actions to perform whenever you restore a backup or upgrade an appliance configured with Hybrid DNS.

Generating a Backup with Hybrid Servers

Whenever you generate an appliance backup, the **Hybrid DNS configuration is automatically retrieved**. For more details regarding backups, refer to the sections [Creating an Instant Backup](#) and [Editing the Backup Settings](#).

Restoring a Backup Containing Hybrid Servers

When you restore an NSD or Unbound server backup, **you must to tick the box "Restore the system configuration"**. Otherwise, BIND service is started and the smart architecture might push an outdated DNS configuration to your physical server instead of your Hybrid configuration. For more details regarding the restoration of a backup, refer to the section [Restoring a Backup File](#).

Chapter 45. DNSSEC

Domain Name System Security Extensions (DNSSEC) is used to strengthen DNS protocol security. It controls the integrity of all DNS answers and ensures that client queries are answered by the proper zone.

By providing origin authentication, it protects the DNS information exchanged between name servers configured with DNSSEC. Within SOLIDserver, it can only be configured on EfficientIP servers and smart architectures managed via SSL, you cannot configure it on other DNS vendors.

DNSSEC ensures data protection from one signed zone to the other, whether the answer is successful or not:

- DNS data in each zone is cryptographically signed** with a couple of public and private Zone Signing Keys (ZSK) that validate the integrity of the data of each zone. As a result:
 - Every RRset of the zone is assigned a new RRSIG record that includes its own signature.
 - The public key is then provided to the resolver or application that validates the integrity of the received RR. The integrity is provided by a chain of trust starting with the public key of a trust anchor.
- NSEC3 records are generated for each RRset**, thus creating an organized chain of all the RRs of the zone that provides an authenticated denial of existence. If the data is supposed to be located in an area of the zone where another RR is located, it means that it does not exist.
- Delegated zones are part of a chain of trust** that ensures that every zone is recognized as legitimate by its parent zone. To implement the security of that relation, each delegated zone ZSK is signed at the parent zone level thanks to a couple of cryptographic Key Signing Keys (KSK) and a Delegation Signer (DS).

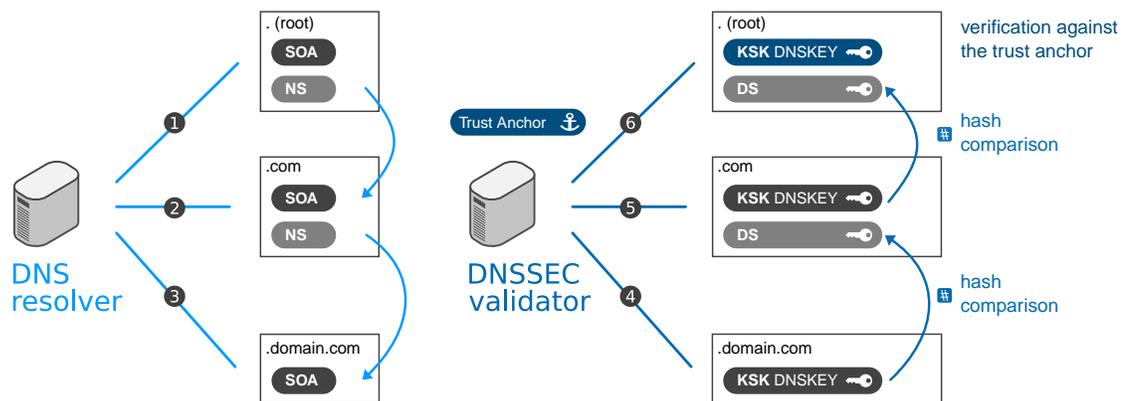


Figure 45.1. DNS resolution and DNSSEC validation

Keep in mind that DNSSEC does not protect whole servers, it only protects the data exchanged between signed zones.

Once DNSSEC is configured, the DNS packages sent and received often exceed 512 bytes, so we recommend configuring EDNS to extend the size of your DNS messages. For more details refer to the section [Configuring EDNS Options at Server Level](#).

Managing DNSSEC on Authoritative Servers

On authoritative servers, implementing DNSSEC means signing at least one of its zones. During the zone signature, a set of signing keys and records are generated:

DNSSEC signing keys

- Two Zone Signing Keys (**ZSK**), to protect the zone data. Every time a signed zone is queried, the ZSK validation process is the following:

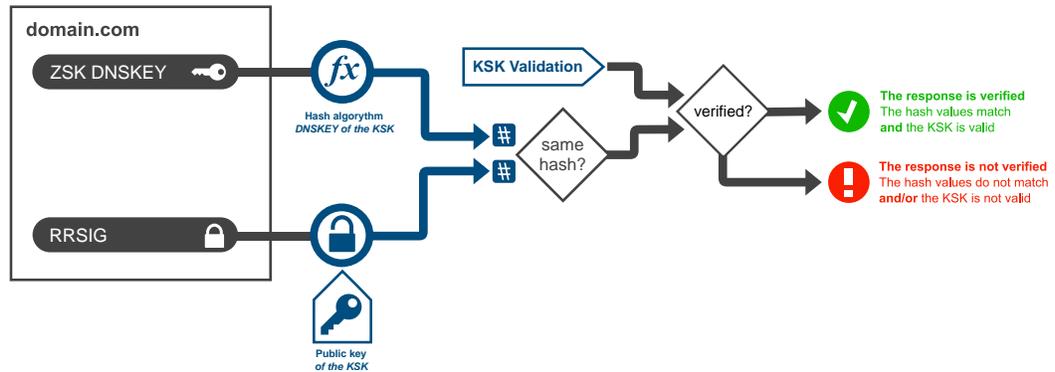


Figure 45.2. The ZSK validation process

One ZSK is active right away and the other is created to replace the first one when it expires or in case of problem.

- One Key Signing Key (**KSK**), to protect the ZSKs. It is active right away. The KSK validation process is the following:

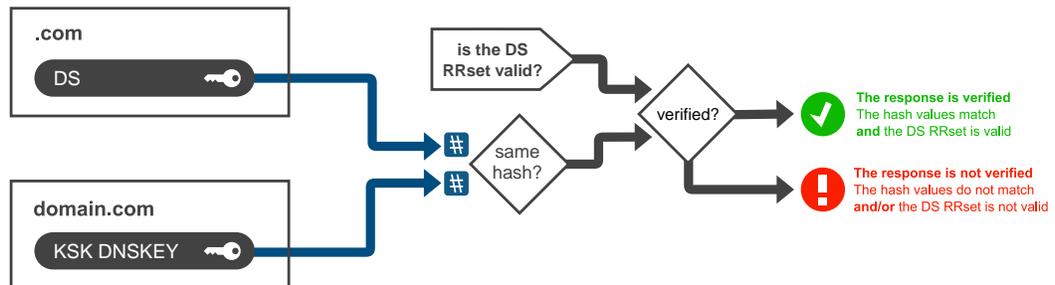


Figure 45.3. The KSK validation process

Once generated, you can manage these keys. For more details, refer to the section [Managing DNSSEC Signing Keys](#).

DNSSEC records

- The **DNSKEY** or Domain Name System KEY. This record contains the public key data of the zone and is used to generate the zone public cryptographic keys, its ZSKs and KSK, it signs and authenticates RRsets.

The DNSKEY is used by DNSSEC clients during the authentication process, it validates the signatures. Three DNSKEY records are generated every time a zone is signed, one for each signing key. The hash of each DNSKEY is compared with the hash of the corresponding RRSIG to ensure that the RRset has not changed.

If you did not sign zones from SOLIDserver, you might need to add DNSKEY records, for more details refer to the section [Adding a DNSKEY Record](#).

In case of unexpected KSK change, you might want to manually add CDNSKEY records to inform the parent zone of these changes. For more details, refer to the [DNSSEC records configuration fields](#) in appendix.

- The **RRSIG** or Resource Record SIGnature. This record stores the digital private signatures, it signs each set of RR of a zone. It does not sign individual records.

The RRSIG guarantees secure DNS operations, its hash is compared with the hash of the DNSKEY to ensure that the RRset has not been changed.

- The **NSEC**, **NSEC3** and **NSEC3PARAM** records. These records significantly extend zone files.

NSEC or Next SECure. This record provides authenticated denial of existence for the records as it points to the next valid host name in the zone for each record. If the requested name is not returned, it does not exist.

NSEC3 or NSEC version 3. This record was designed because NSEC records can help map out the content of the zone. It hashes each label to prevent enumeration.

NSEC3PARAM or Next Secure 3 Parameter. This record assists the authoritative server handling client requests. Thanks to it, they can calculate hashed owner names and choose which set of NSEC3 records are included in the negative responses.

- The **DS**, or Delegation Signer. This record is used to secure delegations between a zone and a subzone. In a parent zone, the DS stores the key tag, algorithm number and a digest of the DNSKEY of its child zone. Both records allow DNSSEC resolvers to authenticate the validity of a subzone. Therefore, once you signed a subzone, you must publish the DS information in the parent zone to make sure it is integrated to the Chain of trust. For more details, refer to the section [Publishing the Delegation Signer in the Parent Zone](#).

In case of unexpected KSK change, you might want to manually add CDS records to inform the parent zone of these changes. For more details, refer to the [DNSSEC records configuration fields](#) in appendix.

Once a zone is signed, querying it means validating its content.

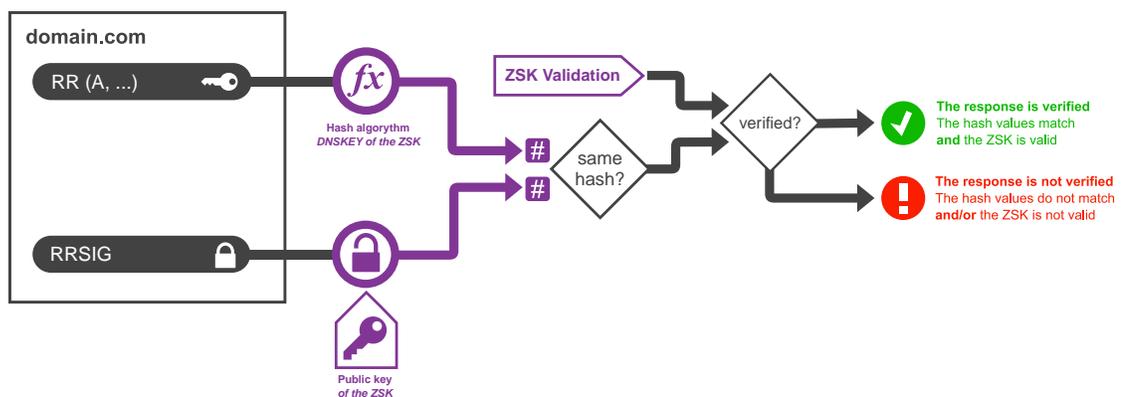


Figure 45.4. RRset validation of a signed zone

Once DNSSEC is configured, the DNS packages sent and received often exceed 512 bytes, so we recommend configuring EDNS to extend the size of your DNS messages. For more details refer to the section [Configuring EDNS Options at Server Level](#).

Signing a Zone

You can sign zones from the page *All zones*. It automatically generates the relevant signing keys and records.

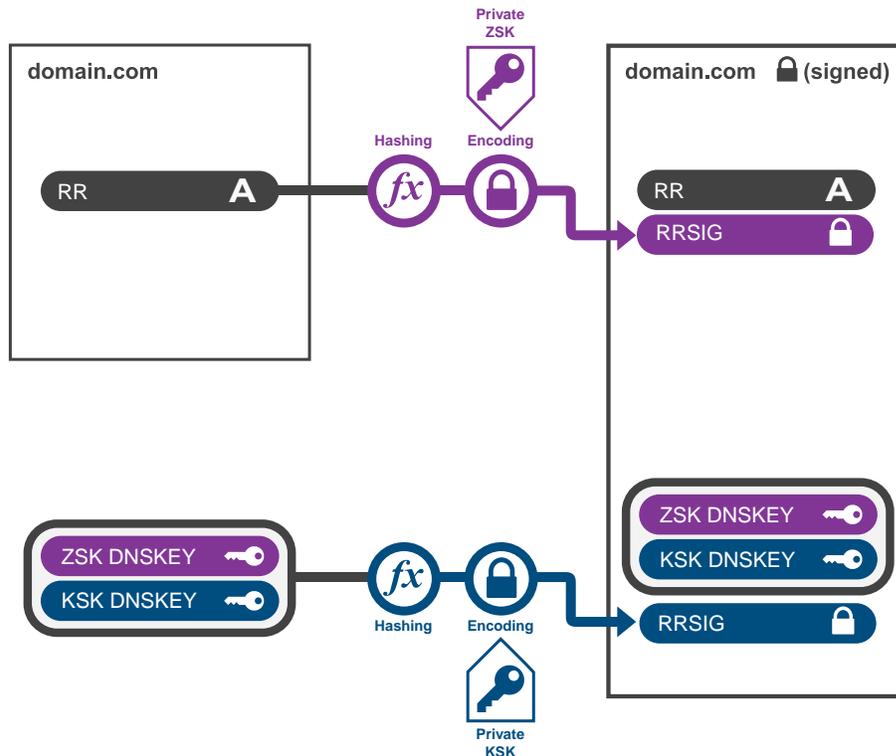


Figure 45.5. The process of signing a zone

Keep in mind that:

- You can only sign master zones managed by smart architectures or EfficientIP DNS servers.
- When you sign a zone for the first time, 3 signing keys are automatically generated, one KSK and two ZSKs. During the zone signature you can configure their life span, a crucial element of the key rollover. For more details, refer to the section [Managing DNSSEC Signing Keys](#).
- When you sign a zone, DNSSEC records are automatically generated once the zone is DNSSEC-compliant. All the records are automatically ordered by the zone and listed on the page *All RRs*, your cannot edit them.
- Once you signed a zone, its signing keys are listed on the page *All DNSSEC keys*. You can use these keys to sign the zone again if you unsigned it or if you want to sign that same zone in a different view of the server or smart architecture.
- We strongly recommend setting up notifications as any problem within a zone can invalidate an entire server.

- If you want to sign zones on a server but also associate it with one or more applications, you must make sure that the A or AAAA record matching the FQDN of each application belongs to a zone that is not signed.

Once you sign a zone that contains records matching any application FQDN, the zone can no longer properly answer signed queries. For more details, refer to the part [Application](#).

- These records are automatically generated once the zone is signed and DNSSEC-compliant, the zone orders the records automatically. They are listed on the page *All RRs*, you cannot edit them.

To sign a zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) you want to sign.
3. In the menu, select **Tools > DNSSEC > Sign zones**. The wizard **Signing zones** opens.
4. Configure the ZSK, either use existing keys or generate a new one.
 - a. To use existing keys, tick the box **Use existing ZSK(s)**. Two lists appear.

Table 45.1. Existing ZSK selection fields

Field	Description
Available ZSK(s)	Select the ZSK of your choice among the <i>Enabled ZSKs</i> generated for the zone and click on  . The key is moved to the list <i>Selected ZSK(s)</i> . You can repeat this operation for as many keys as you need.
Selected ZSK(s)	Lists the ZSKs used to sign the zone. To remove a key from the list, select it and click on  . The key is moved back to the list <i>Available ZSK(s)</i> .

- b. To generate a new ZSK, configure the fields according to your needs.

Table 45.2. New ZSK configuration fields

Field	Description
ZSK - Algorithm	Select the hashing algorithm of the KSK, either <i>RSASHA256</i> , <i>RSASHA512</i> , <i>RSASHA1</i> or <i>DSA</i> . By default, <i>RSASHA1</i> is selected.
ZSK - Encryption	Type in the encryption key size, either <i>1024</i> , <i>2048</i> or <i>4096</i> . By default, <i>1024</i> is specified in the field.
ZSK Validity unit	Select the unit of for the key validity, either <i>Day</i> , <i>Month</i> , <i>Year</i> or <i>Infinity</i> . It applies to the field <i>ZSK - Validity</i> . By default, <i>Month</i> is selected as we recommend setting a ZSK validity of 3 months.
ZSK - Validity	Type in the number of days, months or years that sets the key validity. By default, <i>3</i> is specified in the field as we recommend setting a ZSK validity of 3 months. The field is not displayed if you selected <i>Infinity</i> in the field <i>ZSK Validity unit</i> .

5. Configure the KSK, either use existing keys or generate a new one.
 - a. To use existing keys, tick the box **Use existing KSK(s)**. Two lists appear.

Table 45.3. Existing KSK selection fields

Field	Description
Available KSK(s)	Select the KSK of your choice among the <i>Enabled</i> KSKs generated for the zone and click on + . The key is moved to the list <i>Selected KSK(s)</i> . You can repeat this operation for as many keys as you need.
Selected KSK(s)	Lists the KSKs used to sign the zone. To remove a key from the list, select it and click on - . The key is moved back to the list <i>Available KSK(s)</i> .

- b. To generate a new KSK, configure the fields according to your needs. Keep in mind that your KSK value is probably set by your parent zone.

Table 45.4. New KSK configuration fields

Field	Description
KSK - Algorithm	Select the hashing algorithm of the KSK, either <i>RSASHA256</i> , <i>RSASHA512</i> , <i>RSASHA1</i> or <i>DSA</i> . By default, <i>RSASHA1</i> is selected.
KSK - Encryption	Type in the encryption key size, either <i>1024</i> , <i>2048</i> or <i>4096</i> . By default, <i>2048</i> is specified in the field.
KSK Validity unit	Select the unit of for the key validity, either <i>Day</i> , <i>Month</i> , <i>Year</i> or <i>Infinity</i> . It applies to the field <i>KSK - Validity</i> . By default, <i>Month</i> is selected as we recommend setting a KSK validity of 12 months.
KSK - Validity	Type in the number of days, months or years that sets the key validity. By default, <i>12</i> is specified in the field as we recommend setting a KSK validity of 12 months. The field is not displayed if you selected <i>Infinity</i> in the field <i>KSK Validity unit</i> .

6. Click on **NEXT**. The last page of the wizard opens.
7. Configure the alert notifications, via mail or SNMP trap. You must configure at least one, you can set them both. Make sure the SNMP service is properly configured, for more details refer to the section [Managing the SNMP Service](#).
- a. To set up email alert configuration, make sure the box **Send mail** is ticked. The related fields are displayed.

Table 45.5. Send mail alert configuration parameters

Parameters	Description
Mailing lists	The group of users notified via email. Make sure that you configured an email address for the users of the group you select, otherwise no one in the group will receive the alert.
Additional Mail	Another email address that should receive the alert. Type in an email address and click on ADD , the address is moved to the field <i>Additional Mail List</i> . You can add as many addresses as you want.
Additional Mail List	The list of addresses that receive the alert, they all receive it at the same time no matter the list order. To edit the list, select any entry, it is displayed again in the field. You can either edit the address and click on UPDATE , or remove it from the list by clicking on DELETE . If you made changes that you do not want to save, click on CANCEL .

- b. To set up an **SNMP Trap**, tick the box. The related fields appear. All the fields are compulsory, except the last one.

Table 45.6. SNMP alert configuration parameters

Parameters	Description
SNMP version	The version of SNMP. Select v2c or v3.
SNMP Destination	The IP address of the network management platform. You must type it in.
SNMP Community	The community string that would act as a password to access the SNMP agent. You must type it in.
Raised alert SNMP OID	The custom OID to be sent when the alert is raised. You must type it in. You can use and extend the default OID <i>1.3.6.1.4.1.2440.1.6.1.2.0.1</i> .
Released alert SNMP OID	The custom OID to be sent when the alert is released, this field is optional. If this field is empty, no trap is sent when the alert is released.

8. Click on **OK** to complete the operation. The report wizard opens and closes.

On the page *All zones*, in the column **DNSSEC**, the zone is marked *yes*.

On the page *All DNSSEC keys*, in the column **Status**, the keys generated for the zone(s) are marked *Enabled*. In the column **Life span** you can see which keys are active.

When the zone is signed, it contains DNSSEC records that hold the information of the Trust anchor of the root zone *."*. Therefore, your zone can be included in the chain of trust.

If you signed a subzone, you need to publish the DS record as detailed in the section below.

Publishing the Delegation Signer in the Parent Zone

The Delegation Signer (DS) allows to hold the information of a subzone in a parent zone and ensure DNSSEC authentication at all levels of the chain of trust. Every time you sign a zone, a DS is generated.

Without the DS pointing to the right subzone, DNSSEC resolvers cannot authenticate the subzone as part of the chain of trust and clients cannot access it.

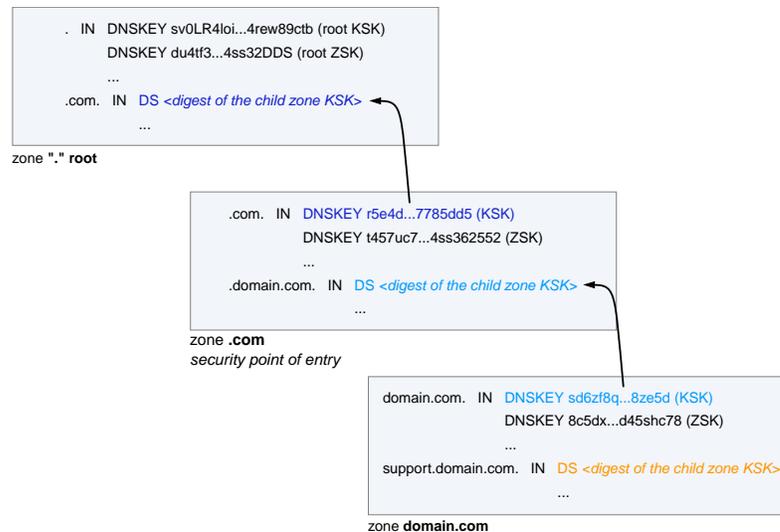


Figure 45.6. The DS record within authoritative zones

All the zones signed with SOLIDserver use the information of the trust anchor of the zone root ".". Therefore, **you need to publish the DS of a subzone if its parent zone was signed**, and only in this case.

Note that if you created a trust anchor for a delegation of private zones, you also need to publish your DS in the signed parent zone. For more details, refer to the section [Adding a Trust Anchor](#).

To publish the DS in the parent zone you must:

1. Retrieve the DS information from the child zone, i.e. copy the DS information on the subzone properties page. Note that you can also export the column DS of a zone in a CSV file, for more details refer to the chapter [Exporting Data](#).
2. If you manage the parent zone, add a DS with that information in the parent zone.

If you do not manage the parent zone of the subzone you signed, copy the entire line of the DS record that suits your needs, paste it in the appropriate file and send it.

Note that:

- You cannot add a DS record in a parent zone if the subzone you are delegating does not contain a DS and a NS record named like the subzone. For more details, refer to the section [Adding an NS Record](#).
- If your parent zone supports it, you can also add a CDS record to let the parent zone retrieve information and be informed of any changes. For more details, refer to the [DNSSEC records configuration fields](#) in appendix.

To publish the DS information of a subzone in its parent zone

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. **Retrieve the DS information**
 - a. At the end of the line of the signed zone of your choice, click on **ⓘ**. The properties page opens.
 - b. In the panel **DS Keys**, click on **▣** to expand it. It contains all the DS records of the zone.

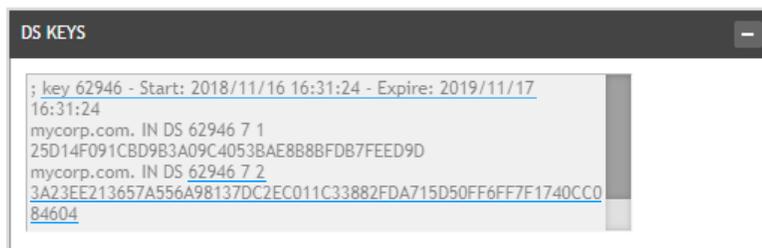


Figure 45.7. Example of a panel DS Keys

For each DS record, you will find:

- On the first line: a *key*, a *Start* date and time and an *Expire* date and time.
- On the next two lines, the DS details. They are displayed twice, in two different digest types. Both lines contain the same zone name and RR type (DS), followed by its complete value.

- c. Identify the DS you need via its *key*. In the image above, we want the details of the key *62946*.
- d. Copy the DS details encrypted with the digest type 2 on the second line. You need all the values displayed right of *DS*, they are separated by a space.

In the image above, the values needed for the zone *mycorp.com* are the **Key Tag** (62946), the **Key Algorithm** (7), the **Digest Type** (2) and the **Digest** (3A23EE213657A556A98137DC2EC011C33882FDA715D50FF6FF7F1740CC084604).

3. Publish the DS information of the subzone in a signed parent zone

If you do not manage the parent zone of the subzone you signed, copy the entire line of the DS record that suits your needs, paste it in the appropriate file and send it.

If you manage the parent zone of the subzone you signed, you need to create a DS record in the parent zone using the child zone DS information:

- a. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
- b. Click on the **Name** of the signed parent zone of your choice. The page **All RRs** opens.
- c. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
- d. In the drop-down list **RR type**, select *DS*.
- e. In the field **RR name**, name the record like the subzone.
- f. In the field or drop-down list **TTL**, you can edit the default value. The field displays the TTL in seconds. Editing the value of the field automatically edits the drop-down list, and vice versa. By default, the TTL is set to *3600* seconds (*1 hour*).
- g. In the field **Key Tag**, paste the child zone key tag.
- h. In the field **Key Algorithm**, paste the child zone key algorithm.
- i. In the field **Digest Type**, paste the child zone digest type.
- j. In the field **Digest**, paste the child zone digest.
- k. Click on **OK** to complete the operation. The report opens and closes. The record is listed on the page **All RRs**, its *Value* contains the *Key Tag*, *Algorithm Key*, *Digest Type* and *Digest* separated by a comma.

Once the parent zone contains a DS record for each child zone, the chain of trust includes both the zone and all its delegated subzones.

Managing DNSSEC Signing Keys

From the page *All DNSSEC keys*, you can manage the ZSKs and KSKs. To manage trust anchors, refer to the section [Managing Trust Anchors](#).

Two **Zone Signing Keys** (ZSK) are generated every time you sign a zone to protect it, one is active right away and the other is ready to replace the first one. You can use existing ZSKs to sign other zones. You can [revoke](#), [disable](#), [enable](#), [delete](#) and [clean](#) unused ZSKs.

One **Key Signing Key** (KSK) is generated every time you sign a zone to protect the ZSKs. You can use existing KSKs to sign other zones. You can [generate](#), [revoke](#), [disable](#), [enable](#), [delete](#) and [clean](#) unused KSKs.

For both key types, you must **monitor and execute the key rollover** whenever necessary. It ensures the security of your DNSSEC zones.

- The **ZSK rollover is automatic** and relies on the two keys generated when you sign a zone. If you suspect a ZSK is compromised you can force an earlier check and generate a new key if relevant. For more details, refer to the section [Executing a ZSK Rollover](#).
- The **KSK rollover is manual**. Only one key is generated when you sign a zone, so you must generate a new one to replace the one protecting the ZSKs of a zone when it expires or if you suspect it is compromised. For more details, refer to the section [Executing a KSK Rollover](#).

Browsing DNSSEC Signing Keys

All the ZSKs, KSKs and trust anchors are listed on the page *All DNSSEC keys*.

To display the list of DNSSEC keys

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

In addition to the keys **Name**, you can display their **Encryption type**, **Key Type** or even **Key tag**.

To display the properties of a DNSSEC key

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. At the end of the line of the key of your choice, click on . The properties page opens.

Executing a ZSK Rollover

The key rollover ensures the security of your DNSSEC system as it makes sure that signed zones are protected by valid keys. The rollover option refreshes the key and generates a new one ready to replace the current one if necessary.

The ZSK rollover is automatic within SOLIDserver for each zone:

- When you sign a zone, two ZSKs are generated and configured with the same parameters. The first one is published in the zone and immediately active, the second one is published but inactive until the active ZSK expires.
- When the active ZSK expires, the second one is activated and a third key is published and inactive until the second ZSK expires, and so forth. All ZSKs share the parameters of the signed zone.

The whole set of active keys is checked daily.

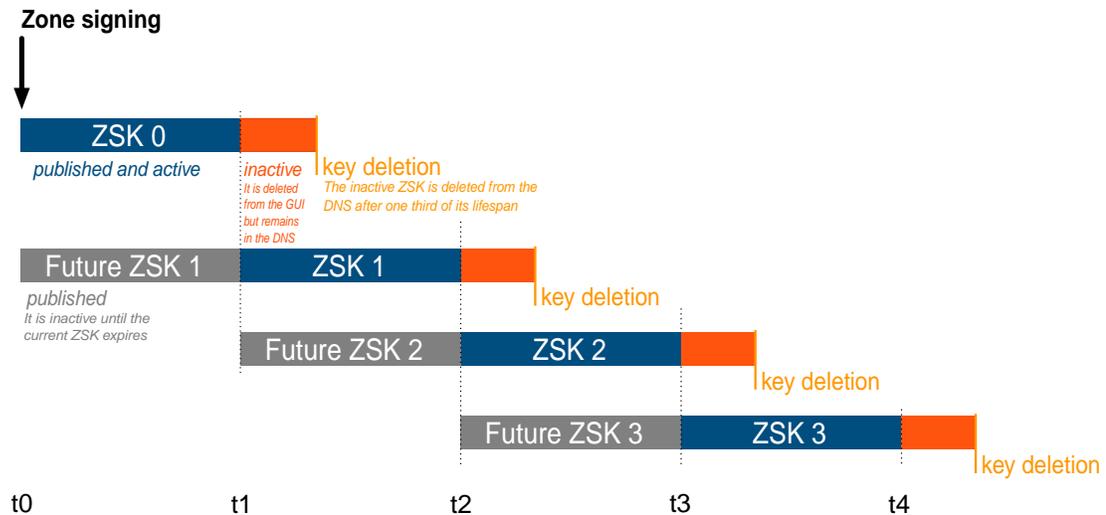


Figure 45.8. ZSK rollover

All keys are listed on the page *All DNSSEC keys*. Inactive ZSKs are automatically deleted from the GUI, but remain in the DNS for one third of their lifespan before being deleted from the DNS as well. These inactive keys are kept in the DNS because that information has been cached by other servers. Therefore, the information of the inactive ZSK is still used to secure the records, unless you edited them after the activation of the new ZSK. Once the ZSK is deleted from the DNS, all the records are automatically signed again using the currently active ZSK.

Even if the ZSK rollover is automatic on signed zones, you can force an earlier check if you suspect they are compromised.

The option *Force automatic ZSK rollover* allows to make sure that the selected ZSK is not compromised or soon to expire. If it is the case, a new one is generated to replace it.

If a key is compromised you must revoke it, for more details refer to the section [Revoking a ZSK](#).

To execute the ZSK rollover

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. In the menu, select **Tools > Force automatic ZSK rollover**. The wizard **Forcing the refresh of all the ZSKs** opens.
5. Click on **[OK]** to complete the operation. The report opens and closes. A new ZSK is generated only if necessary.

Executing a KSK Rollover

The key rollover ensures the security of your DNSSEC system as it makes sure that signed zones are protected by valid keys. The rollover option refreshes the key and generates a new one ready to replace the current one if necessary.

The KSK rollover is manual, so the option *KSK rollover* is all the more useful.

Ideally the KSK has to be replaced once a year.

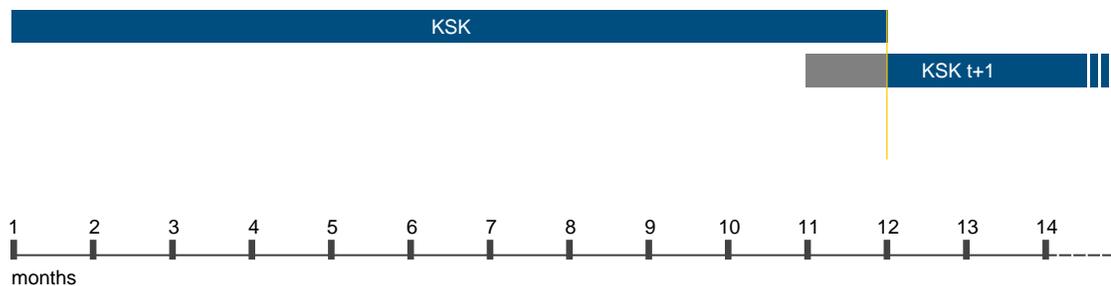


Figure 45.9. KSK rollover

If a KSK is compromised, you must revoke it. For more details, refer to the section [Revoking a KSK](#).

To execute the KSK rollover

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. Tick the KSK of your choice.
5. In the menu, select **Edit > KSK rollover**. The wizard **Performing a KSK rollover** opens.
6. Click on **OK** to complete the operation. The report opens and closes. A new KSK is generated only if necessary.

Generating a New KSK

Some time before the KSK expires, you are notified via email alert and/or an SNMP trap depending on what you configured when you signed the zone. You may also need a new KSK to safely revoke or disable the current one.

To properly enable the rollover of the KSK and protect the ZSKs of your zone you must:

1. Generate a new KSK.
2. Publish the Delegation Signer (DS) that signs the new KSK in your parent zone.

We strongly recommend setting an alert to be notified when the KSK you generate is about to expire, using the column *Time left* for instance. For more details, refer to the section [Managing Alerts](#).

To generate a new KSK

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. **Generate the KSK**
 - a. Tick the signed zone(s) of your choice.
 - b. In the menu, select **Tools > DNSSEC > Generate new KSK**. The wizard **Generating a new KSK** opens.
 - c. Configure the KSK according to your needs:

Table 45.7. New KSK configuration fields

Field	Description
KSK - Algorithm	Select the hashing algorithm of the KSK, either <i>RSASHA256</i> , <i>RSASHA512</i> , <i>RSASHA1</i> or <i>DSA</i> . By default, <i>RSASHA1</i> is selected.
KSK - Encryption	Type in the encryption key size, either <i>1024</i> , <i>2048</i> or <i>4096</i> . By default, <i>2048</i> is specified in the field.
KSK Validity unit	Select the unit of for the key validity, either <i>Day</i> , <i>Month</i> , <i>Year</i> or <i>Infinity</i> . It applies to the field <i>KSK - Validity</i> . By default, <i>Month</i> is selected as we recommend setting a KSK validity of 12 months.
KSK - Validity	Type in the number of days, months or years that sets the key validity. By default, <i>12</i> is specified in the field as we recommend setting a KSK validity of 12 months. The field is not displayed if you selected <i>Infinity</i> in the field <i>KSK Validity unit</i> .

d. Click on **OK** to complete the operation. The report wizard opens and closes.

3. Publish the DS of the new KSK in the parent zone

- a. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
- b. Click on **All DNSSEC keys**. The page opens.
- c. At the end of the line of the new KSK, click on **⌵**. The properties page opens.
- d. In the panel **DS**, copy the content of the list *DS*.
- e. Transmit the DS to the parent zone. For more details, refer to the procedure [To publish the DS information of a subzone in its parent zone](#).

Some parent zones may also require the DNSKEY record of the new KSK.

Revoking a KSK

If a KSK is compromised you should revoke it. Revoking a KSK invalidates its corresponding DNSKEY record for the zone and can protect it from attacks.

To properly revoke a KSK you must:

1. Generate a new KSK and notify the parent zone to make sure you do not invalidate the whole zone by revoking the compromised key. For more details, refer to the section [Generating a New KSK](#).
2. Revoke the compromised KSK.

Note that **you cannot revoke a KSK if it is the only one protecting the ZSKs of a zone**, it would break the chain of trust and prevent from successfully querying your zone records, the zone would be invalidated.

To revoke a KSK

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.

4. Generate a new KSK

- a. Make sure that you have another KSK, a valid one, protecting the zone and that you notified the parent zone.
- b. If you do not have two KSKs attached to the zone listed on the page or if you have not notified the parent zone of the new KSK generation, refer to section [Generating a New KSK](#).

5. Revoke the compromised KSK

- a. Tick the KSK of your choice.
- b. In the menu, select  **Edit > Revoke KSK Keys**. The wizard **Revoking Key Signing Keys** opens.
- c. Click on to complete the operation. In the column **Type**, the KSK is now marked *KSK (invalidated)* but its **Status** is still *Enabled*.

The KSK that now protects the zone is edited:

- Its **Start** date now matches the time and date of the revocation.
- Its **Validity** and **Time left** evolve to protect the zone until the intended *End* time and date of the KSK you revoked.

Revoking a ZSK

If a ZSK is compromised, you must revoke it. As there are two ZSKs protecting a signed zone, revoking a ZSK triggers the following behaviors:

- **If the active ZSK is revoked**

1. The revoked ZSK is deleted from the GUI and the DNS database.
2. The published ZSK is activated to replace it. The new ZSK lasts longer than the initial time configured when you signed the zone: it lasts for the remaining time of the revoked key plus its own lifespan.
3. Another ZSK is generated and published. It lasts as long as the active ZSK. When the active ZSK expires, it replaces it, this time it lasts for the initial period configured when you signed the zone.

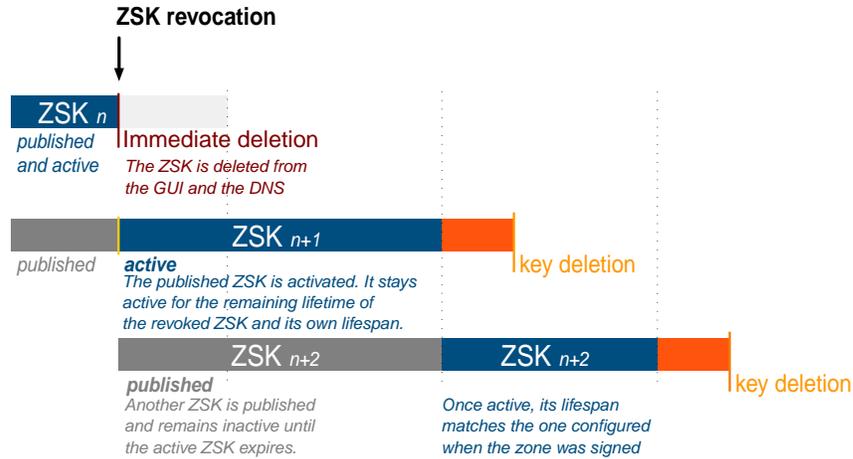


Figure 45.10. The mechanism when you revoke the active ZSK

- **If the published ZSK is revoked**

1. The revoked ZSK is deleted from the GUI and the DNS database.
2. Another ZSK is generated and published, matches the active ZSK configuration. It is activated when the active ZSK expires.

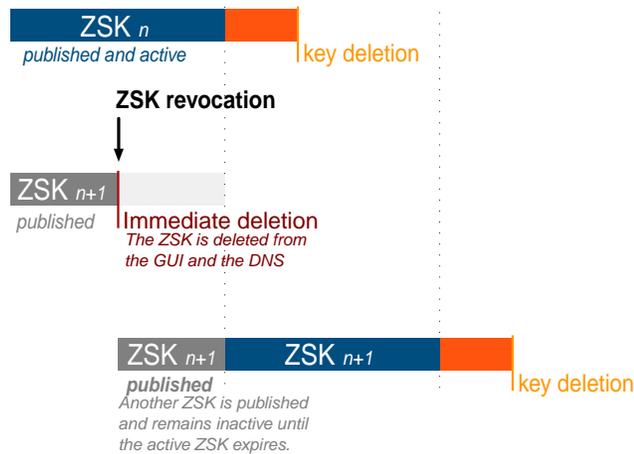


Figure 45.11. The mechanism when you revoke the published ZSK

To revoke a ZSK

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. Tick the ZSK(s) of your choice.
5. In the menu, select **Edit > Revoke ZSKs**. The wizard **Revoking ZSKs** opens.
6. Click on **OK** to complete the operation. The wizard closes and the page refreshes.

If you revoked an inactive key, it is replaced by a new key with a different **Key tag** and the same **Start** and **End** dates.

If you revoked an active key, it is replaced by the inactive key and a new one is generated to replace it when it expires.

Disabling Signing Keys

You can disable ZSKs and KSKs, this may be useful if you need to complete your DNSSEC configuration and then push it again on the server.

By default the generated ZSKs and KSKs are enabled, this does not mean that they are currently used to protect the zone. In the column **Life span**, you can see that the KSK and one ZSK is active, inactive keys are marked *Not started yet*.

Keep in mind that:

- Disabling a key does not delete it. To delete a key, refer to the section [Deleting Signing Keys](#).
- You can disable keys whether they are currently active or not.
- Disabling a key deletes the corresponding DNSKEY and RRSIG records in the relevant zone.
- The active ZSK and the KSK both protect the zone. If you disable the KSK of a zone but not its ZSKs, or vice versa, the zone appears *Broken* on the page *All zones*, in the column *DNSSEC*.
- When you disable a ZSK:
 - If the ZSK is active, the zone is no longer protected. The inactive ZSK does not take over to protect the zone, disabling an active key cancels the rollover. The zone is no longer protected and marked *Broken* until you enable the key again.
 - If the ZSK is inactive, the zone is protected until the active ZSK expires. The disabled ZSK does not take over when the active ZSK expires, the zone is *Broken*.
- When you disable a KSK:
 - If the KSK is active and no other KSK is configured for the zone, it breaks the chain of trust and can affect a zone and its subzones.
 - If the KSK is active and another KSK is configured for the zone, both keys are protecting your ZSKs. If you disable one, the other KSK is still protecting your zone.
 - If the KSK is inactive but the zone is still protected by the active one, the zone is protected until the active KSK expires. The zone is *Broken* until you replace the KSK.
- Disabling a key prevents from using it to sign a zone. Disabled keys are not listed in the fields *Available ZSK(s)* and *Available KSK(s)* of the wizard *Signing a zone*.
- It is possible to select all the keys of a zone and disable them but it is not recommended. Indeed, disabling all keys at once partially unsigns the zone(s) as the DNSKEY and RRSIG records are deleted but it does not properly purges all the DNSSEC records. To properly unsign zones, refer to the section [Unsigning a Zone](#).

To avoid breaking the chain of trust before disabling a KSK or ZSK you should:

- **Make sure the key is inactive**, and not currently used to sign any zone.
- **Make sure another key is ready to protect the zone**. Two ZSKs are generated when you sign a zone, however you might need to generate a new KSK for the zone(s), before disabling a KSK. For more details, refer to the section [Generating a New KSK](#).

To disable a DNSSEC key

This operation should only be performed by administrators with good knowledge of DNSSEC mechanisms.

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. From the column **Status**, you can filter the list to only display *Enabled* or *Disabled* keys.
5. Filter the list to only list enabled keys.
6. Tick the ZSK(s) or KSK(s) of your choice.
7. In the menu, click on **✎ Edit > Enable**. The wizard opens.
8. Click on **OK** to complete the operation. In the column **Status**, the selected key(s) is marked **Disabled**.

Enabling Signing Keys

You can enable again a disabled ZSK or KSK. Keep in mind that:

- Enabling a key recreates the corresponding DNSKEY and RRSIG records in the relevant zone.
- Enabling a key again is not instantaneous. Until the cache of all the servers that queried the zone expires, the new key information, including its life span, is not up-to-date and the change is not taken into account.
- If the key you disabled was active, enabling it again allows to use it to protect the zone. Inactive keys enabled again only change status, they are used when the active key expires.
- If the zone was *Broken*, enabling again the proper keys changes its *DNSSEC* status to *yes* on the page *All zones*.
- Enabled keys are available again in the list of existing keys of the zone signature wizard.

To enable DNSSEC keys

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. From the column **Status**, you can filter the list to only display *Enabled* or *Disabled* keys.
5. Filter the list to only list disabled keys.
6. Tick the ZSK(s) or KSK(s) of your choice.
7. In the menu, click on **✎ Edit > Enable**. The wizard opens.
8. Click on **OK** to complete the operation. In the column **Status**, the selected key(s) is marked **Enabled**.

Deleting Signing Keys

You can only delete a KSK or ZSK if it has expired, is *Disabled* or was generated for a zone that was unsigned.

To delete a signing key

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. Filter the list to only display unused keys.
5. Tick the key(s) of your choice.
6. In the menu, click on **Delete**. The wizard **Delete** opens.
7. Click on **OK** to complete the operation. The key is no longer listed.

Cleaning Unused Signing Keys

By default, the unused ZSKs and KSKs are automatically deleted from the page *All DNSSEC keys* every night. However, at any point you can remove unused keys yourself.

The option *Clean unused keys* looks for expired signing keys in all relevant zones and deletes them if relevant.

To clean unused keys

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. In the menu, select **Tools > Clean unused keys**. The wizard **Remove unused DNSSEC keys** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The relevant keys are deleted.

Unsigning a Zone

You can unsign zones to stop using DNSSEC to secure your DNS organization.

Unsigning a zone disables DNSSEC for a zone. Note that disabling DNSSEC on a zone is different than disabling the signing keys of a zone. For more details, refer to the section [Disabling Signing Keys](#).

Before unsigning a zone, keep in mind that:

- Unsigning a zone disables its signing keys and purges all the DNSSEC records generated during the signature. The zone is no longer considered secure by other signed zones.
- Unsigning a parent zone breaks the chain of trust set with its child zone(s). As only the records generated automatically are purged, the DS record you published for any child zone is still part of the parent zone.

You must delete the DS record published in the parent zone for each child zone to ensure it cannot compromise it. Once the parent zone is unsigned, its DS records are no longer authenticated or secure.

- Unsigning a zone signed using HSM must be unsigned using the HSM as well. The server managing the zone(s) must have the box *Enable HSM* ticked when you unsign the zones to

ensure that both SOLIDserver and the HSM module are notified that one zone is no longer part of the DNSSEC delegation. For more details, refer to the chapter [HSM](#).

Once a zone is unsigned, its signing keys are still listed on the page *All DNSSEC keys* and can be used to sign the zone again.

To unsigned DNSSEC zones

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. Tick the zone(s) of your choice.
3. In the menu, select **Tools > DNSSEC > Unsign zones**. The wizard **Unsigning zones** opens.
4. Click on **OK** to complete the operation. The report wizard opens and closes.

On the page *All zones*, in the column **DNSSEC**, the zone is marked *no*.

On the page *All DNSSEC keys*, in the column **Status**, the keys are marked *Unused*, if they were only used with the zone you unsigned.

5. If you unsigned a parent zone, you must delete the DS record of each of its signed zones as it is no longer signed, secure and authenticated. For more details, refer to the section [Deleting Resource Records](#).

Managing DNSSEC on Recursive Servers

On recursive servers, DNSSEC relies on resolvers that you include to a chain of trust using trust anchors.

Within SOLIDserver, you can enable DNSSEC validation, manage the trust anchors and disable the validation.

Enabling DNSSEC Validation

Just like the DNS, DNSSEC validation relies on resolvers. They must be part of a chain of trust.

You can set Efficient DNS servers and smart architectures as DNSSEC resolvers and associate them with a trust anchor.

The chain of trust ensures that clients are directed to valid zones. All queries and answers are signed and compared at every DNS lookup node to authenticate the exchange and make sure that both sides can be trusted. That way, at all relevant levels, a verified encrypted signature provides validating resolvers with the correct path to secured zones and prevents directing clients toward *bogus* IP addresses. This validation can be set all the way up to the TLD.

The starting point of the chain of trust is the trust anchor. Once configured on a validating resolver, it allows the resolver to validate the integrity of the records sent by DNS clients and ensure the chain of trust between a zone and its parent. Therefore, a zone or subzone has to be secured before being linked to the secured zone it is a delegated from. The trusted anchor of the parent zone then covers the secured child zones that are delegated from it.

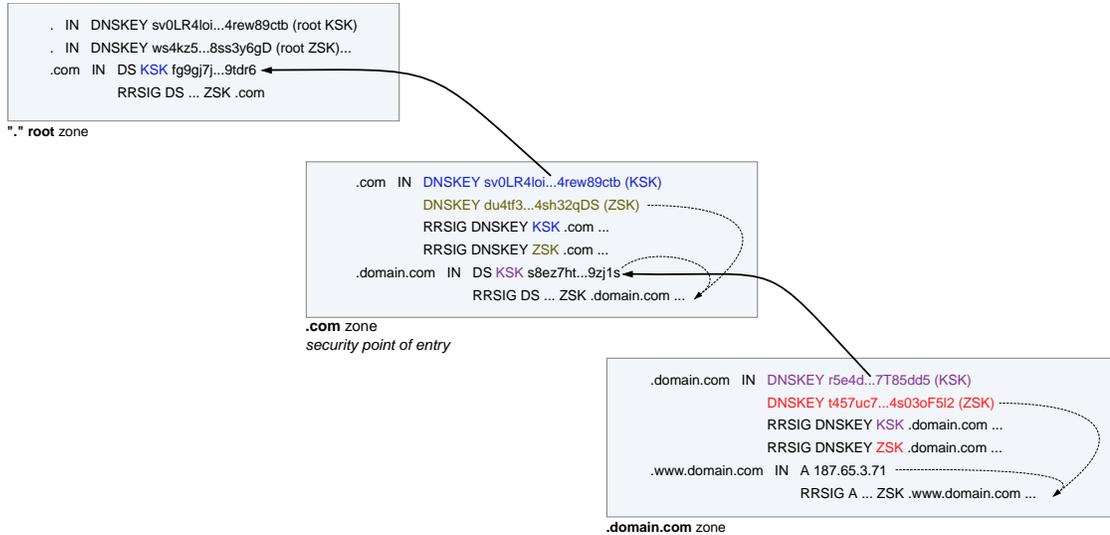


Figure 45.12. The DNSSEC chain of trust from root to subzones

Keep in mind that:

- On smart architectures, enabling DNSSEC validation makes all the physical servers they contain DNSSEC compliant.
- Maintaining a valid chain of trust is paramount because broken chains of trust would result in data being marked as *bogus* and may cause entire zones or subzones to become invisible to verifying clients. To verifying clients, a secure zone is either part of a chain or trust or insecure.
- By default, ICANN trust anchors are available on the page *All DNSSEC keys*. If you want to add another trust anchor, refer to the section [Adding a Trust Anchor](#).
- DNSSEC validation is all the more efficient if your chain of trust is complete before setting a server as resolver. Indeed, when a DNSSEC resolver receives a response from an unsigned zone that has a signed parent:
 - It confirms with the parent that the zone was intentionally left unsigned. It verifies, via signed and validated NSEC/NSEC3 records, that the parent zone contains no DS records for the child.
 - If the DNSSEC resolver can prove that the zone is secure, it accepts the response.
 - If the DNSSEC resolver cannot prove that the zone is secure, it assumes that the response is insecure and probably a forgery, rejects the response and logs an error.
- You can enable a server as DNSSEC resolver without signing the zones that the server manages. This ensures that the information sent out to its DNS clients is valid and protected.
- A DNS servers used as DNSSEC resolver cannot validate records matching nodes in the module Application. Once an application is deployed on a GSLB enabled physical server, all its nodes replace DNS answers and do not update RRSIG records. For more details on nodes, refer to the part [Application](#).

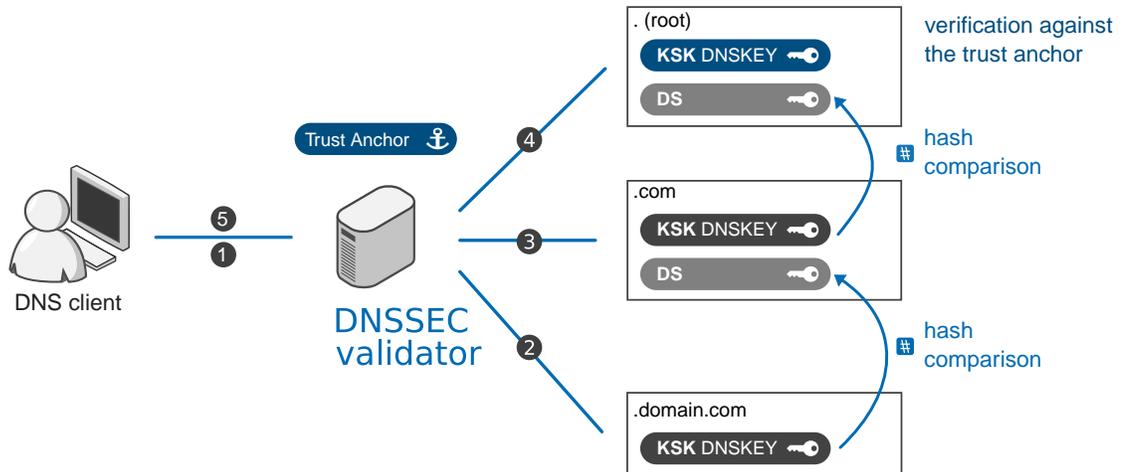


Figure 45.13. DNSSEC validation

To enable DNSSEC validation

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on **⚙**. The properties page opens.
3. In the panel **DNSSEC**, click on **EDIT**. The wizard **Edit DNSSEC properties** opens.

The box is not available for servers managed via a smart architecture, in this case to tick the box you must edit the smart architecture.

4. Tick the box **Use DNS as DNSSEC resolver**. The wizard refreshes.
5. In the list **Available Trust Anchor**, select a trust anchor and click on **+**. The trust anchor is moved to the list **Configured Trust Anchors**.

To remove a trust anchor from the list, select it and click on **-**. It is moved back to the list *Available Trust Anchor*.

6. Click on **OK** to complete the operation. The wizard closes. In the panel **DNSSEC**, the **DNSSEC resolution** is now *Enabled* and the list **Trust Anchors** contains the chosen trust anchor(s).

On the page *All servers*, in the column **DNSSEC** the server is now marked *Yes*.

On the properties page of the trust anchor, in the panel *DNS servers using this trust anchor*, the server is listed.

Once DNSSEC is configured, the DNS packages sent and received often exceed 512 bytes, so we recommend configuring EDNS to extend the size of your DNS messages. For more details refer to the section [Configuring EDNS Options at Server Level](#).

Managing Trust Anchors

From the page *All DNSSEC keys*, you can manage the trust anchors used to configure resolvers for DNSSEC validation.

You can [add trust anchors](#), [use one trust anchor on several resolvers](#) and [delete](#) them.

Browsing the Trust Anchor Database

To display the list of trust anchors

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. At the end of the line of the key of your choice, click on **■**. The properties page opens.

All trust anchors have a unique **Name** on the page *All DNSSEC keys*. By default, ICANN trust anchors are available on the page.

Adding a Trust Anchor

If you want to set up your own chain of trust from one of your signed zones, you can add a trust anchor.

If a local trust anchor is available, it is used to verify the zone without comparing the KSK protecting the zone against all its signed parent zone.

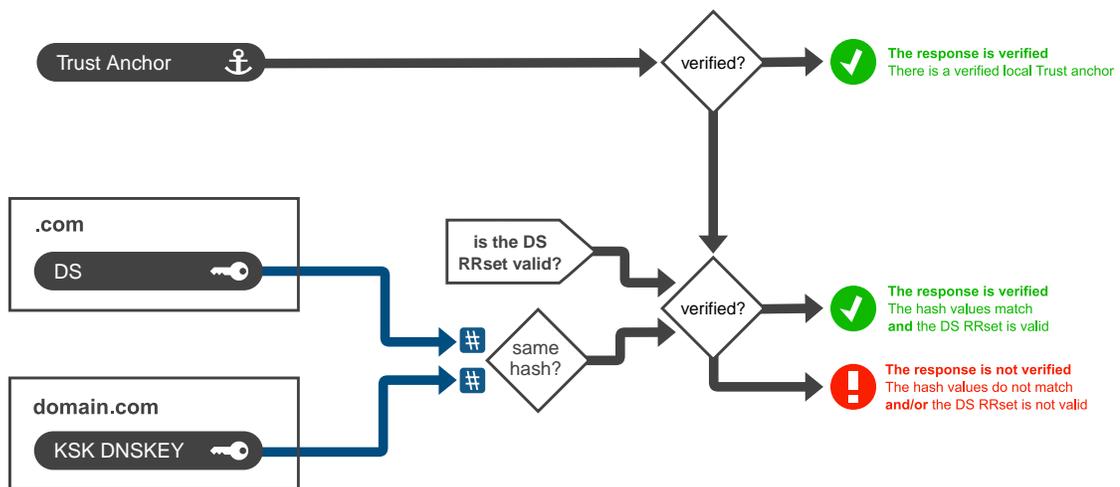


Figure 45.14. The KSK is verified if a local trust anchor is available

To successfully set up your chain of trust you must:

1. Sign the zone used at the top of the chain of trust.
2. Retrieve the trust anchor information on the properties page of the KSK of that zone.
3. Add the new trust anchor to the page *All DNSSEC keys*.
4. Use the trust anchor on the relevant DNSSEC resolver.

To add a DNSSEC trust anchor

1. In the sidebar, go to **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on **»** to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.

4. **Retrieve the trust anchor information**
 - a. At the end of the line of the KSK of the zone of your choice, click on . The properties page opens.
 - b. In the panel **Trust anchor**, copy the content of the list *Key*.
 - c. In the breadcrumb, click on **DNSSEC Key** to go back to the page *All DNSSEC keys*.
5. **Add the trust anchor**
 - a. In the menu, click on **+ Add**. The wizard **Add trust anchor** opens.
 - b. In the field **Key**, paste the trust anchor information you just retrieved.
 - c. Click on  to complete the operation. The trust anchor is now listed on the page *All DNSSEC keys*. In the column **Zone**, the name of the zone it was retrieved from is listed.
6. Now you must add the trust anchor to the resolver managing your zones as detailed in the procedure [To enable DNSSEC validation](#).

Using One Trust Anchor On Several Resolvers

You can use one trust anchor on several resolvers, provided that they support it.

To use one trust anchor on several resolvers

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on  to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. At the end of the line of the trust anchor of your choice, click on . The properties page opens.
5. In the panel **DNS servers using this Trust Anchor**, click on . The wizard **DNSSEC chain of trust configuration** opens.
6. In the **DNS server list**, select a server and click on . The server is moved in the list **Selected**. Repeat this action for as many servers as needed.

To stop using a trust anchor on a particular server, select it in the list **Selected** and click on . It is moved back to the **DNS server list**.

7. Click on  to complete the operation. The servers are now listed in the panel **DNS servers using this Trust Anchor**.

Deleting a Trust Anchor

You can only delete a trust anchor no longer used by any DNSSEC resolver.

To delete a DNSSEC trust anchor

1. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
2. In the breadcrumb on the right of **All zones**, click on  to display additional pages.
3. Click on **All DNSSEC keys**. The page opens.
4. Filter the list to only display unused keys.

5. Tick the trust anchor(s) of your choice.
6. In the menu, click on  **Delete**. The wizard **Delete** opens.
7. Click on to complete the operation. The key is no longer listed.

Disabling DNSSEC Validation

You can disable DNSSEC validation on a resolver, you must edit it and untick the relevant box.

Once the validation is disabled, the server is no longer part of the chain of trust of the trust anchor. Both the server and the trust anchor information are updated.

To disable DNSSEC validation

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on . The properties page opens.
3. In the panel **DNSSEC**, click on . The wizard **Edit DNSSEC properties** opens.
4. Untick the box **Use DNS as DNSSEC resolver**. The wizard refreshes and no longer displays the lists *Available Trust Anchor* and *Configured Trust Anchors*.

The box is not available for servers managed via a smart architecture, in this case to untick the box you must edit the smart architecture.

5. Click on to complete the operation. The wizard closes. In the panel *DNSSEC*, the **DNSSEC resolution** is now *Disabled* and the list **Trust Anchors** is empty.

On the page *All servers*, in the column **DNSSEC** the server is now marked *No*.

On the properties page of the trust anchor, in the panel *DNS servers using this trust anchor* no longer lists the server.

Chapter 46. HSM

The goal of this chapter is to detail the procedures to sign Master zones with DNSSEC using a Hardware Security Module (HSM). This operation is only possible with an EfficientIP DNS server, either managed by a smart architecture or not.

A Hardware Security Module (HSM) is a physical crypto processor dedicated to generating, storing and managing encryption keys. It provides strong authentication, through accelerated cryptographic operations, and multiple levels of security. With SOLIDserver, you can sign EfficientIP DNS servers' zones with DNSSEC using HSM encryption.

DNSSEC is an extension to DNS that provides a way to authenticate DNS responses, thus preventing DNS traffic interference such as man in the middle attacks. For DNSSEC to function properly, it is essential for the private keys, ZSK and KSK, to be protected¹. With HSM, the private keys never leave the module thus preventing any exposition to potential attacks, on the network or host computer. SOLIDserver is compatible with the Thales nShield 500 HSM.

To ensure a secure communication with a DNS server, the HSM relies on:

- **The Security World (SW), a closed architecture that describes the global cryptographic environment**

The SW is used to encrypt recovery data and authorize the creation of Operator Card Sets (OCS) that can be required for the module to function. For redundancy, it can manage several *HSM servers*, i.e. modules, but is supervised using a unique Administrator Card Set (ACS).

You can create your own SW or integrate your SOLIDserver appliance to an existing one.

- **The Remote File System (RFS) that contains the keys and configuration data**

Every key is generated on the RFS, sent to the module for encryption and sent back to the RFS for storage. Any time the server needs a key, the encrypted key is transmitted by the RFS to the module(s) for decryption and sent straight to the server.

There should be one RFS by Security World stored either locally on your SOLIDserver appliance or, if you want to integrate your SOLIDserver appliance to an existing SW, on an external server².

Prerequisites

- Only users of the group *admin* can perform the operations detailed in this chapter.
- Only the Thales nShield Connect 500 HSM is supported.
- Your license must be valid and include the DNS module.
- Your SOLIDserver appliance should be added to the HSM authorized client list using the proper ACS. For more details, refer to the documentation of your HSM.
- Master zones to be signed with the HSM must be managed by an EfficientIP DNS server, whether in a smart architecture or not. This feature is not compatible with other types of DNS servers.

¹For more details, refer to the chapter [DNSSEC](#).

²The RFS storage being remote in essence, you need to choose a server that offers enough data security to store the RFS outside of the HSM.

- To properly secure appliances in High Availability, it is strongly advised to set up the HSM on each of them before enrolling the Hot Standby.

Limitations

- The HSM slot is common to all the HSM queries. You cannot set it per DNS server.
- Deleting private keys from SOLIDserver does not trigger automatic deletion on the RFS. For more details regarding how to delete them from the RFS, refer to the documentation of your HSM.
- A card from the OCS needs to be inserted in the HSM for encryption to work. Softcards, i.e. virtual security tokens, are not supported.
- K/N quorum is not supported, only 1/N is. This means you cannot require more than one card to use a module.
- When the HSM feature is enabled on an EfficientIP DNS server, all the zones you sign with it are signed using the HSM, i.e. you cannot sign some zones with the module and others without it.
- All the Master zones signed using the HSM must also be unsigned using the HSM.

Browsing HSM Servers

When the HSM has been fully configured, the page **Services configuration** provides you with the properties you have set for the HSM module(s). Once enrolled and enabled, the status of the service *DNS server (HSM)* should be  **OK**. The service *DNS server (Named)* should be  **Disabled**.

To display the HSM servers list

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, the different properties of the HSM are listed below the service *nFast Hardserver*.

The screenshot shows the 'Services configuration' page in SOLIDserver. The page title is 'Services configuration' and the user is 'sadmin'. The page displays a list of 28 services with columns for Name, Configuration, Description, Running, Enabled, and Status. The services listed include SSH server, Accounts admin, SFTP - SCP - RSYNC, Accounts xfer, TFTP server, Upload Authorization: Disabled, Mail - SMTP, Default source mail, Alert source mail, SMTP relay, Outgoing mail server, HTTP server, HTTP webservice, SSL Certificate, DNS server (named), DNS server (PSM), nFast Hardserver, HSM protection : ocs, RFS server, HSM servers, HSM-EP, DHCP server, DHCPv6 server, and SNMP server.

Name	Configuration	Description	Running	Enabled	Status
SSH server			Started	Enabled	OK
Accounts admin					OK
SFTP - SCP - RSYNC					OK
Accounts xfer					OK
TFTP server			Stopped	Disabled	Disabled
Upload Authorization: Disabled					OK
Mail - SMTP					OK
Default source mail	noreply@efficio.com				OK
Alert source mail	noreply@efficio.com				OK
SMTP relay					OK
Outgoing mail server	None				OK
HTTP server			Stopped	Disabled	Disabled
10.55.18.27/09/2016 (FORCE UPDATE)					OK
HTTP webservice					OK
SSL Certificate	Apache SSL Cert Base				OK
DNS server (named)			Stopped	Disabled	Disabled
DNS server (PSM)			Started	Enabled	OK
nFast Hardserver			Started	Enabled	OK
HSM protection : ocs	PIN:****				OK
RFS server	10.10.17.84				OK
HSM servers					OK
HSM-EP: 10.0.240.250	ESN:8C80-D0CB-175F-KNETI+HASH:9f6c27eeb189a876eddc010b-805be70724af68	nShield			OK
DHCP server			Started	Enabled	OK
DHCPv6 server			Started	Enabled	OK
SNMP server			Started	Enabled	OK
SNMP community: public	Read only; Restriction: default				OK

Figure 46.1. HSM service configuration

HSM Protection

On the page **Services configuration**, in the column *Name*, the service **HSM Protection** is followed either by **OCS**, if the module is configured to function only when a physical smart card is inserted, or **Module** if no card is required.

If a PIN code has been specified for the OCS, the column *Configuration* displays the mention **PIN code:******.

RFS Server

On the page **Services configuration**, the column *Configuration* of the service **RFS server** indicates the IP address of the server on which the RFS is stored. If the RFS is stored locally on the SOLIDserver appliance, the column displays the mention *Local* instead.

HSM Servers

For redundancy, you can manage several modules, also referred as *HSM servers* in SOLIDserver, within the same Security World. They are listed under the service **HSM servers** on the page **Services configuration**:

- The column *Name* indicates the name and IP address of the module(s) enrolled by SOLIDserver.
- The column *Configuration* displays the electronic serial number (ESN) and hash of the KNETI key (KNETI HASH) in use for each module.
- The column *Status* provides information regarding the modules - or *HSM servers* - you manage.

Table 46.1. HSM server statuses

Status	Description
OK	The HSM module or service is operational.
Missing RFS	The RFS is missing and the HSM server is not operational.
Not identified	The HSM module is not identified and therefore not operational.
Not enrolled	The HSM module is not enrolled and therefore not operational.

In case a module fails, whether you are managing only one or several modules, there is no visual indication on the GUI.

However, the logs in the file *named.log* indicate that SOLIDserver has switched back from the service *DNS server (HSM)* to *DNS server (named)* to sign zones with DNSSEC. For more details regarding how to display logs, refer to the section [Syslog](#).

```

❑ 06/10/2016 11:45:08 client 10.0.200.15#2053 updating zone 'only1.hsm/IN' RRSIG/NSEC/NSEC3 update failed sign failure
❑ 06/10/2016 11:45:08 /mnt/proj/package-6-0-0/nessy/bind-9.10/lib/dns/pkcs11rsa_link.c 468 pkcs_C_SignFinal Error = 0x00000030
❑ 06/10/2016 11:45:08 client 10.0.200.15#2053 updating zone 'only1.hsm/IN' adding an RR at 'sig6.only1.hsm' A 6.2.2.2

```

Figure 46.2. HSM error log

In both cases, it is necessary to stop the service **DNS server (HSM)** and restart it before you can use it again to sign zones. For more details regarding how to restart a service, refer to the section [Starting or Stopping a Service](#).

Setting Up the HSM

Setting up the HSM for DNSSEC zone signature requires to:

1. Set up the service *nFast Hardserver* that represents the Security World. This implies configuring the **HSM protection**, the **RFS server** and the IP address of each physical module, called **HSM servers** in the GUI. For more details, refer to the section [Setting Up the Service nFast Hardserver](#).
2. Switch from the classic service *DNS Server (named)* to *DNS Server (HSM)*. For more details, refer to the section [Switching from DNS Server \(named\) to DNS Server \(HSM\)](#).
3. If you plan on using HSM with High Availability, set up the service *nFast Hardserver* and switch to the proper DNS service as detailed in the the two previous steps on both appliances **before** you enroll the Hot Standby. For more details, refer to the chapter [Centralized Management](#).

You can however set up your HSM on appliances already in HA and force the synchronization within a minute. For more details, refer to the section [Enable HSM on Appliances Already in High Availability](#).

4. Enable HSM on the servers managing zones that you want to sign using HSM. For more details, refer to the section [Signing Zones Using HSM and DNSSEC](#).

Setting Up the Service nFast Hardserver

To enroll an HSM with SOLIDserver, you need to enable and configure the service *nFast Hardserver* from the Administration module.

If you want to specify a communication slot different from the default ones, follow the procedure in the section [Adding a Custom Communication Slot](#) below. Otherwise, go straight to the section [Enabling the Service nFast Hardserver](#).

Adding a Custom Communication Slot

To communicate with the Thales nShield Connect 500 HSM, SOLIDserver uses the preconfigured slots *492971158* for OCS protection and *492971157* for Module protection.

If you want to specify a different slot than the default ones, you need to add a new item in the Registry database. Otherwise, go straight to the section [Enabling the Service nFast Hardserver](#).

To add a custom slot ID

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the menu, click on  **Add**. The wizard **Registry database Add an item** opens.
4. In the field **Name**, type in *module.dns.hsm_slot* .
5. In the field **Value**, type in the number of your slot.
6. Click on to complete the operation. The report opens and closes. The page **Registry database** is visible again.
7. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
8. In the section **System**, click on **Services configuration**. The page **Services configuration** is visible again.

Enabling the Service nFast Hardserver

To properly enroll a module with SOLIDserver, you need to enable the service *nFast Hardserver* before configuring it.

If you plan on using HSM with High Availability, make sure to set up and enable the service *nFast Hardserver* on both the future Master and the Hot Standby before configuring the association.

To enable the service nFast Hardserver

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, find the service **nFast Hardserver**.
4. In the column **Enabled**, click on **Disabled** to enable the service. The wizard **Enable a service** opens.
5. Click on to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.
6. In the menu, select  **Tools** > **Apply configuration**. The wizard **Commit the system configuration changes** opens.
7. Click on to complete the operation. The report opens and closes. In the column **Status**, the **nFast Hardserver** is marked  **OK**.

Configuring the Service nFast Hardserver

The service *nFast Hardserver* allows to enroll one or several modules. To configure the service, you can:

- Either specify an existing Security World by declaring the IP address of its Remote File System,
- Or use your SOLIDserver as the local RFS, in which case, it would contain the Security World needed to authenticate your data exchanges.

To configure the service nFast Hardserver

Only users of the group *admin* can perform this operation.

1. On the HSM, add your SOLIDserver to the authorized clients list. For more details, refer to the documentation of your HSM.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
4. In the column **Name**, click on the service *nFast Hardserver*. The wizard **HSM server configuration** opens.
5. In the drop-down list **Protection method**, set the configuration that suits your needs:

Table 46.2. HSM protection configuration

Field	Description
Module	The module doesn't need a physical smart card from the Operator Card Set (OCS) to function. By default, this option is selected.
OCS	The module functions only when a card from the OCS is inserted. If selected, the page refreshes and displays the field <i>OCS PIN code</i> .
OCS PIN code	Type in a password for the OCS, it should be the same for every card of a given set. If you want to integrate SOLIDserver to an existing Security World, use the related PIN code.

6. Click on **NEXT**. The next page of the wizard opens.
7. Set the configuration that suits your needs for the module you are enrolling:

Table 46.3. HSM protection configuration

Field	Description
Name	Type in a name for the module.
IP address	Type in the IP address of the module.
Description	Type in a description for the module.
Override default parameters	Tick this box to set additional parameters. If selected, the page refreshes and displays the boxes <i>Force enrollment</i> and <i>Privilege</i> .
Force enrollment	Tick this box to force the reconfiguration of the module if it has once been enrolled by SOLIDserver.
Privilege	Tick this box if you want SOLIDserver to request a connection with privileges to the module.

8. Click on **ADD** to commit your addition. The page refreshes and the server is moved to the list **HSM servers**.
 - To update an entry, select the module of your choice, edit the data and click on **UPDATE**.
 - To delete an entry, select the module of your choice and click on **DELETE**.
 - To discard the latest changes, click on **CANCEL**.
 - To add another module to the Security World, repeat steps 7 and 8.
9. Click on **NEXT**. The next page of the wizard opens.
10. In the drop-down list **RFS**, set the configuration that suits your needs:

Table 46.4. RFS storage configuration

Field	Description
Local	The HSM uses SOLIDserver as the Remote File System to store the Security World for authentication and data exchanges. By default, this option is selected.
Remote	The HSM uses an external server to store the RFS, which can be the one of an existing Security World. The page refreshes and displays the field <i>IP</i> .
IP	Set the IP address of the external server used to store the RFS.

11. Tick the box **Expert mode** to set additional parameters:

Table 46.5. RFS expert mode configuration

Field	Description
Force setup	Tick this box to overwrite the previous RFS server configuration.
Address	Type in the IP address of a client you want to grant write access to the RFS.
Name	Type in a name for the client.
Description	Type in a description for the client.

12. If you specified a client with write access, click on **ADD** to commit your addition. The page refreshes and the client is moved to the list **Allowed clients**.
- To update an entry, select the client of your choice, edit the data and click on **UPDATE**.
 - To delete an entry, select the client of your choice and click on **DELETE**.
 - To discard the latest changes, click on **CANCEL**.
 - To grant another client write access to the RFS, repeat steps 11 and 12.
13. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again. The enrolled modules are listed under the line *HSM servers* along with their ESN and KNETI-HASH. The statuses of the *nFast Hardserver* and module(s) should be  **OK**. For more details, refer to the section [Browsing HSM Servers](#).

Switching from DNS Server (named) to DNS Server (HSM)

To enable the DNS service dedicated to HSM, once the service *nFast Hardserver* is configured, you must:

1. Disable the service *DNS server (named)*.
2. Enable the service *DNS server (HSM)*.

These instances cannot be running at the same time. To stop using the HSM, refer to the section [Best Practices to Stop Using the HSM](#).

Note that, if you enable the service *DNS server (HSM)* before adding any HSM server, the service Status should be  **No HSM found**. In this case, you need to add the HSM server and restart the service *DNS server (HSM)*.

To switch from DNS Server (named) to DNS Server (HSM)

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, find the service **DNS server (named)**.
4. Disable the service **DNS server (named)**:
 - a. In the column **Enabled**, click on **Enabled** to disable the service. The wizard **Disable a service** opens.
 - b. Click on **OK** to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.
 - c. In the menu, select **Tools > Apply configuration**. The wizard **Commit the system configuration changes** opens.
 - d. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again and the service is marked **Disabled**.
5. In the column **Name**, find the service **DNS server (HSM)**.
6. Enable the service **DNS server (HSM)**:
 - a. In the column **Enabled**, click on **Disabled** to enable the service. The wizard **Enable a service** opens.
 - b. Click on **OK** to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.
 - c. In the menu, select **Tools > Apply configuration**. The wizard **Commit the system configuration changes** opens.
 - d. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again and the service is marked **Enabled**.

When the service *DNS server (HSM)* is enabled, you can start signing zones through EfficientIP DNS servers.

Enable HSM on Appliances Already in High Availability

It is strongly advised to set up the HSM before configuring High Availability. However, you can still force HSM synchronization by [Setting Up the Service nFast Hardserver](#) and [Switching from DNS Server \(named\) to DNS Server \(HSM\)](#) on both Master and Hot Standby appliances and then following the procedure below.

To force HSM synchronization with the Hot Standby

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** appliance.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
4. In the column **Name**, type in *module.system.file_sync_period* and hit Enter. The key is the only one listed.
5. In the column **Value**, click on the value. The wizard **Registry database edit a value** opens.

6. In the field **Value**, type in *60* to force the replication on the Hot Standby within the minute. The default value is *3600*.
7. Click on to complete the operation. The report opens and closes. The page **Registry database** is visible again.
8. Wait for more than 60 seconds, in order for the Hot Standby to synchronize.
9. Connect to the **Hot Standby**.
10. Restart the DNS and HSM services:
 - a. Disable the Service nFast Hardserver. For more details, refer to the section [Disabling the Service nFast Hardserver](#).
 - b. Switch from *DNS Server (HSM)* to *DNS Server (named)*. For more details, refer to the section [Switching from DNS Server \(HSM\) to DNS Server \(named\)](#).
 - c. Enable the Service nFast Hardserver. For more details, refer to the section [Enabling the Service nFast Hardserver](#).
 - d. Switch from *DNS Server (named)* to *DNS Server (HSM)*. For more details, refer to the section [Switching from DNS Server \(named\) to DNS Server \(HSM\)](#).
11. Repeat steps 1 to 5.
12. In the field **Value**, type in the value *3600*.
13. Click on to complete the operation. The report opens and closes. The page **Registry database** is visible again.

Signing Zones using DNSSEC and HSM

To sign zones with DNSSEC and your HSM, you need to tick the box **Enable HSM** on an EfficientIP DNS server and, if need be, the smart architecture that manages it. This feature is not compatible with other types of DNS servers. Once the feature is enabled, signing zones follows the same procedure as detailed in the chapter [DNSSEC](#), but the zone data is encrypted and stored in the HSM.

Keep in mind that:

- **From a server with the HSM feature enabled, all the Master zones are signed using the HSM.**

You cannot sign some zones with the HSM and others without it.

- **All the Master zones signed using the HSM must also be unsigned using the HSM.**

When you unsign the zones, you must make sure that the box **Enable HSM** is still ticked.

To enable HSM on an EfficientIP server or a smart architecture

Only users of the group *admin* can perform this operation.

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The properties page opens.
3. In the panel **Main properties**, click on . The wizard **Edit a DNS server** opens.
4. Click on until the wizard displays the box **Enable HSM**.

5. Tick the box **Enable HSM** to use HSM when you sign the server's zones with DNSSEC.
6. Click on **OK** to complete the operation. The report opens and closes. The server or smart architecture is listed and uses HSM to authenticate DNSSEC keys.
7. Repeat steps 3 to 8 for as many servers or smart architectures as needed.

Best Practices to Stop Using the HSM

At any time, you can completely stop using encrypted data communication via the module. To do so, you need to connect to SOLIDserver GUI using the credentials of a user belonging to the group *admin* and follow the steps below:

1. Delete all the *HSM servers* from the list as detailed in the section [Deleting an HSM Server](#).
2. Disable the service *nFast Hardserver* as detailed in the section [Disabling the Service nFast Hardserver](#).
3. Switch back from the service *DNS server (HSM)* to *DNS server (named)* as detailed in the section [Switching from DNS Server \(HSM\) to DNS Server \(named\)](#).
4. Disable HSM on the EfficientIP DNS servers and smart architectures on which it is enabled as detailed in the section [Disabling HSM on DNS Servers and Smart Architectures](#).

Deleting an HSM Server

At any time you can delete an HSM server from the list. This action purges all the configuration files of the module thus preventing any communication with SOLIDserver. It also forces a cleanup of all the related files. The HSM server should no longer be listed on the page *Services configuration*.

To delete an HSM server

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, click on the service *nFast Hardserver*. The wizard **HSM server configuration** opens.
4. Click on **NEXT**. The next page of the wizard opens.
5. In the list **HSM servers**, click on the module you want to delete. The page refreshes and the fields display the module's details.
6. Click on the button **DELETE**. The page refreshes, the fields are emptied and the server is removed from the list.
7. Repeat steps 5 and 6 for as many servers as needed.
8. Click on **NEXT**. The next page of the wizard opens.
9. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again and the HSM server is no longer listed.

Disabling the Service nFast Hardserver

Once the HSM servers are deleted, you must disable the service *nFast Hardserver*.

To disable the service nFast Hardserver

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, locate the service **nFast Hardserver**.
4. In the column **Enabled**, click on **Enabled** to disable the service. The wizard **Disable a service** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.
6. In the menu, select **Tools > Apply configuration**. The wizard **Commit the system configuration changes** opens.
7. Click on **OK** to complete the operation. The report opens and closes. The page refreshes, the service **nFast Hardserver** is marked as **Disabled**.

Switching from DNS Server (HSM) to DNS Server (named)

As the service *DNS server (HSM)* is not fully functional at this point, it is necessary to enable back the service *DNS server (named)* to handle the requests resolution.

To switch from DNS Server (HSM) to DNS Server (named)

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
3. In the column **Name**, find the service **DNS server (HSM)**.
4. Disable the service **DNS server (HSM)**:
 - a. In the column **Enabled**, click on **Enabled** to disable the service. The wizard **Disable a service** opens.
 - b. Click on **OK** to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.
 - c. In the menu, select **Tools > Apply configuration**. The wizard **Commit the system configuration changes** opens.
 - d. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again and the service is marked **Disabled**.
5. In the column **Name**, find the service **DNS server (named)**.
6. Enable the service **DNS server (named)**:
 - a. In the column **Enabled**, click on **Disabled** to enable the service. The wizard **Enable a service** opens.
 - b. Click on **OK** to complete the operation. The report opens and closes. The service starts automatically but is listed in red because the configuration is still pending.

- c. In the menu, select **✖**. **Tools > Apply configuration**. The wizard **Commit the system configuration changes** opens.
- d. Click on **OK** to complete the operation. The report opens and closes. The page **Services configuration** is visible again and the service is marked **Enabled**.

Disabling HSM on DNS Servers and Smart Architectures

To remove the remaining link between the HSM and the DNS server(s) you used to sign zones, you need to untick the box **Enable HSM**.

To disable HSM on an EfficientIP server or a smart architecture

Only users of the group *admin* can perform this operation.

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on **⚙**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a DNS server** opens.
4. Click on **NEXT** until the wizard displays the box **Enable HSM**.
5. Untick the box **Enable HSM** to stop using HSM when you sign the server's zones with DNSSEC.
6. Click on **OK** to complete the operation. The report opens and closes. The server or smart architecture is listed and does not use HSM anymore to authenticate DNSSEC keys.
7. Repeat steps 3 to 9 for as many servers or smart architectures as needed.

Chapter 47. Monitoring and Reporting DNS Data

SOLIDserver provides a set of tools to monitor DNS servers and generate reports.

- The **alerts** that you can set on the DNS pages allow to customize your monitoring. For more details, refer to the chapter [Managing Alerts](#).
- A set of **statistics** are available in dedicated panels of the properties page of DNS servers, as detailed in the section [Monitoring DNS Servers From their Properties Page](#).
- A set of data sampling **analytics** are available on the page *Analytics*, as detailed in the section [Monitoring DNS Servers From the Page Statistics](#).
- A couple of tools allow to monitor a DNS server **querylog** and **answerlog**, as detailed in the section [Monitoring DNS Queries and Answers](#).
- A number of **reports** on servers, views and zones are available, as detailed in the section [Generating DNS Reports](#).

Monitoring DNS Servers From their Properties Page

On the properties page of a physical or smart DNS server, some panels are dedicated to monitoring queries and changes. The monitoring panels are the following:

DNS server statistics <server-name>

Displays query dedicated charts for the physical servers, except DNS Hybrid servers.

- *Success/Recursion* shows the number of successful answers and the number of answers that were not cached, in queries per second.
- *Failure/NXRRSET/NXDOMAIN/Referral/Duplicate/Dropped* shows the answering behavior of the server, in queries per second.
- *Authoritative/Recursive* shows authoritative and recursive answers, in queries per second.
- *DNS firewall (RPZ)* shows answers matching a RPZ rule configured on the server, in queries per second . For more details about Response Policy Zone, refer to the chapter [DNS Firewall \(RPZ\)](#).

This data is retrieved using SNMP, therefore, the graphs are empty if the SNMP is not configured properly. To edit the SNMP parameters of an EfficientIP DNS server, refer to the section [Editing the SNMP Monitoring Parameters of an EfficientIP DNS Server](#).

State log

Displays the server logs.

Audit

Displays all the latest changes performed on the server by the user logged in.

Note that you might have additional panels called **Guardian - <data>** on the properties page, if you have enabled DNS Guardian on your appliance. For more details, refer to the part [Guardian](#).

Keep in mind that you can zoom in and out of the charts or decide the period and data to display. For more details refer to the section [Charts](#).

To display a DNS physical server query statistics

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on . The properties page open.
3. To open the panel **DNS server statistics <physical-server-name>**, click on .

Note that this panel can be displayed on any dashboard, like the other gadgets. For more details, refer to the section [Assigning a Gadget from a Resource Properties Page](#).

To display a DNS server state log

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on . The properties page open.
3. To open the panel **State log**, click on . The panel content retrieves the server state in the logs: *OK*, *KO*, *Invalid settings...* and the time and date for each.

To display a DNS server audit

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on . The properties page open.
3. To open the panel **Audit**, click on . The panel displays the latest changes in the database.

The *Date* and time it occurred, the *Service* used, the *User* performing the operation and the server basic information: *DNS name*, *DNS type* and *Architecture* if relevant.

By default, it lists the changes carried out by the user logged in, but if they belong to a group with access to the changes from all users, the panel displays all the operations ever performed. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).

Monitoring DNS Servers From the Page Analytics

The page **Analytics** provides data sampling via DNS query dedicated *Top 50* lists for the EfficientIP DNS physical servers and the Hybrid DNS servers you manage. There are some limitations detailed in the section [DNS Analytics Limitations](#).

The analytics functionality is enabled by default and samples the DNS traffic over specific periods of time. By default, it offers 5-minute samples. If you want to set a shorter or bigger periodicity, refer to the section [Configuring the Analytics Retrieval](#).

You can set up an alert on the entries displayed. For more details, refer to the chapter [Managing Alerts](#).

Accessing the Page Analytics

The page *Analytics* offers dedicated *Top 50* data samples based on the DNS queries of the physical servers.

The *Top 50* information compares data retrieved over a specific periodicity, a limited period of time, that is set by default to a sample time of 5 minutes. You can edit the sampling period following the procedure in the section [Configuring the Periodicity](#).

Note that the page might contain *Top 50* samples and *DNS Guardian* data. If you display the *Analytics* of a specific server, either DNS data or DNS Guardian data is displayed. For more details regarding DNS Guardian data, refer to the section [Monitoring Guardian Statistics from the GUI](#).

To display the page Analytics

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
3. Click on **Analytics**. The page refreshes.

To display the page Analytics of a specific physical server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on **■**. The properties page opens.
3. In the breadcrumb on the right of **All servers**, click on **»** to display additional pages.
4. Click on **Analytics**. The page refreshes.

Each column provides and compares message information over the specified sample time:

Table 47.1. The columns on the page Analytics

Column	Description
Server	The name of the physical server, it is display the page for all the servers. Click on a server name to display only the analytics of that server.
Start date	The time and date when the data retrieval started.
End date	The time and date when the data retrieval stopped. The end date respects the number of minutes set in the periodicity.
Period	The periodicity set for the sample of queries. It is set on a physical server properties page. For more details, refer to the section Configuring the Periodicity .
Domain	The name of the domain queried by the client <example.com>.
TLD	The TLD of the domain queried by the client, <com>. This column is only displayed for <i>Top 50 TLDs / Source IP</i> and <i>Top 50 Source IPs</i> .
Source IP	The IP address of the client querying the information. This column is only displayed for <i>Top 50 TLDs / Source IP</i> , <i>Top 50 Destination IPs</i> and <i>Top 50 TLDs / Source IPs</i> .
Destination IP	The IP address of the client receiving the information. This column is only displayed for <i>Top 50 Destination IPs</i> .
OpCode	The operation codes returned for each query during the selected period: <i>Query</i> , <i>Status</i> , <i>Notify</i> , <i>Update</i> ... This column is only displayed for <i>Opcodes</i> .
RCode	The return code sent with each query response during the selected period: <i>Noerror</i> , <i>Servfail</i> , <i>Nxdomain</i> ... This column is only displayed for <i>Rcodes</i> .
Query type	The DNS resource record returned to provide the answer to the client: <i>SOA</i> , <i>NS</i> , <i>A</i> , <i>AAAA</i> , <i>PTR</i> , <i>MX</i> ... This column is only displayed for <i>Query types</i> .
Total queries	The total number of messages on the DNS traffic during the selected period, it includes the queries and the codes returned.
Number of hits	The exact number of times the domain/TLD/record/code was queried or returned for the selected period. For instance, the number of times <example.com> was queried by the IP

Column	Description
	address <1.2.3.4> during a 5-minute period. For the <i>Top 50 Domains / source IP</i> and the <i>Top 50 TLDs / source IP</i> , the number of hits is based on the value of the columns <i>Source IP</i> and <i>End date</i> ; e.g., the number of times an A record was queried by the IP address <1.2.3.4> during a 5-minute period.
Ratio	The percentage of <i>Number of hits</i> compared with the <i>Total queries</i> for the selected period. For instance, the percentage that represents the A records queried to find the IP address of <example.com> compared with all the messages on the traffic for a period of 5 minutes.

Note that the columns **Start date** and **End date** can be filter using the keyword **last** to display the data retrieved in the last *X* minutes, *X* being the periodicity set for the server. If no data is retrieved in that period, the list is empty.

Displaying the DNS Analytics

From the page *Analytics*, you can display the physical server data samples retrieved from the messages. It focuses by default on a sample period of 5 minutes to draw *Top 50* comparisons.

To display specific DNS analytics data

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the column **Name**, click on the server or smart architecture of your choice to display the zones it contains.
3. In the breadcrumb on the right of the server name, click on **»** to display additional pages.
4. Click on **Analytics**. The page refreshes.
5. Under the menu, in the drop-down list **Display**, select the data of your choice. The page refreshes, the selected data is displayed.

All available analytics are detailed in the table below. If the drop-down list only contains data named *Guardian - <sample>*, refer to the section [Displaying Guardian Analytics Tops](#).

6. Under the drop-down list, you can tick the box **Automatic refresh** to automatically refresh the data listed every minute. To edit the page refresh frequency, refer to the section [Editing the Automatic Refresh Frequency](#).

You can edit the sample time following the procedure in the section [Configuring the Periodicity](#).

Table 47.2. DNS analytics

Statistic	Description
Top 50 Domains	The top 50 queried domains on the DNS traffic during the configured <i>Period</i> , identified using their name and TLD.
Top 50 Domains / source IP	The top 50 domains queried by specific IP addresses on the DNS traffic during the configured <i>Period</i> , identified using the domain name and TLD and the querying client IP address.
Top 50 Destination IPs	The top 50 answered clients on the DNS traffic during the configured <i>Period</i> , identified using their IP address.
Top 50 Source IPs	The top 50 querying clients on the DNS traffic during the configured <i>Period</i> , identified using their IP address.
Top 50 TLDs	The top 50 TLDs queried on the DNS traffic during the configured <i>Period</i> , identified using the domain TLD.
Top 50 TLDs / source IP	The top 50 TLDs queried by specific IP addresses on the DNS traffic during the configured <i>Period</i> , identified using the domain TLD and the querying client IP address.

Statistic	Description
Opcodes	All the operation codes on the DNS traffic during the configured <i>Period</i> , identified using their name.
Rcodes	All the return codes on the DNS traffic during the configured <i>Period</i> , identified using their name.
Query types	All the queries answered on the DNS traffic, identified using the RR type.

Configuring the Analytics Retrieval

You can configure the analytics retrieval according to your needs. You can edit the sampling period, or *periodicity*, the page automatic refresh frequency, the data retrieval frequency and even its purge frequency.

Editing the Automatic Refresh Frequency

On the page **Analytics**, the box **Automatic refresh** allows to automatically refresh the page display every 60 seconds. You can edit this frequency via a registry database key.

To edit the analytics automatic refresh frequency

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *www.dns.stat.refresh* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value. The wizard registry database edit a value opens.
5. In the field **Value**, type in the number of seconds of your choice. By default, it is set to *60*.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new key value is visible.

Configuring the Periodicity

By default, the data sampling compare the DNS traffic over a sample periodicity of 5 minutes. The sample time specified in the column *Period* on the page *Analytics*.

You can configure a shorter or larger periodicity on each physical server individually.

Note that no matter the periodicity, the data is available on the page at a frequency specified through the rule 380. To edit that rule, refer to the next section.

To edit the DNS analytics periodicity of a physical server

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the physical server of your choice, click on . The properties page opens.
3. Open the panel **DNS analytics**. It indicates if the retrieval is enabled and the periodicity.
4. Click on . The wizard **Configure DNS analytics** opens.
5. In the drop-down list **Periodicity (min.)**, select the period of your choice: *1*, *5*, *10* or *15* minutes. By default, *5* is selected.
6. Click on to complete the operation. The page refreshes, the properties is visible again.

Configuring the DNS Analytics Retrieval Frequency

The frequency to which the analytics are displayed in the GUI is set by the rule 380, *Retrieval of the DNS server analytics data*. By default, every 5 minutes it displays the data comparison results for messages sampled during 1, 5, 10 or 15 minutes, depending on the configured periodicity.

No matter the periodicity you set on the physical server, the data is available in the GUI depending on the rule configuration.

To edit the rule 380 that sets the DNS analytics data retrieval

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #**, type in *380* and hit Enter. The rule is the only one listed.
4. At the end of the line, click on . The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a rule** opens.
6. Click on **NEXT**. The page **Rule filters** opens.
7. Edit the drop-down lists **Day(s) of the week**, **Date of the month**, **Month**, **Hour** and/or **Minute** according to your needs. By default, only the drop-down list *Minute* is set to *Every 5 minutes*.
8. Click on **OK** to complete the operation. The page refreshes, the properties is visible again.

Configuring the DNS Analytics Purge Frequency

You can configure the purge mechanism of the analytics retrieval. By default, it is based on:

- The data age. The rule 382, *Configuration of the DNS analytics purge*, deletes data older than 30 days. You can set it to delete data earlier or later.
- A lines count. A registry key deletes data if the analytics database exceeds 100,000 lines - each *Top 50*, *Opcodes*, *Rcodes* and *Query types* can reach that many lines. You can set a lower or higher threshold.

Both thresholds work together: once the number of days or the number of lines is met, the unwanted data the deleted.

No matter the way you want to purge your database, **keep in mind that if you set very high thresholds, you may slow down your appliance** because the database contains too much information.

To edit the rule 382 that purges the DNS analytics

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #**, type in *382* and hit Enter. The rule is the only one listed.
4. At the end of the line, click on . The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a rule** opens.
6. Click on **NEXT**. The page **Rule filters** opens.
7. Edit the drop-down lists according to your needs. By default the rule is executed daily at 23:30.

Table 47.3. Filters of the rule 382

Column	Default value
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, 23 is selected.
Minute	Select a period of time, minutes-wise. By default, 30 is selected.

8. Click on . The page **Rule parameters** opens.
9. In the field **Number of days**, type in the number of days above which you want the logs database to be purged. By default, it is set to 30: logs older than thirty days are automatically deleted.
10. Click on to complete the operation. The page refreshes, the properties is visible again.

To create a threshold to purge DNS analytics

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the menu, click on **Add**. The wizard **Registry database Add an item** opens.
4. In the field **Name**, type in *dns.stats.limit*.
5. In the field **Value**, type in the number of lines above which the data is purged. The default value is *100,000*.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the key is listed.

Exporting the Analytics

From the page *Analytics*, you can export the data listed in a CSV, HTML, XML, XLS or PDF file.

Like any other export, you can retrieve the data immediately or schedule it. For more details, refer to the section [Configuring Exports](#).

Disabling the Analytics

At any time, you can stop retrieving the analytics for any physical server.

To disable the DNS analytics retrieval on a physical server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server of your choice, click on . The properties page opens.
3. Open the panel **DNS analytics**. It indicates if the retrieval is enabled and the periodicity.
4. Click on . The wizard **Configure DNS analytics** opens.
5. Untick the box **Enable analytics collection**. The page refreshes, the drop-down list *Periodicity (min.)* is no longer visible.

- Click on to complete the operation. The page refreshes, the properties is visible again. In the panel, the field **Enable analytics collection** is marked **no**.

DNS Analytics Limitations

- The analytics are only available for EfficientIP and Hybrid DNS servers.
- The analytics data is only retrieved based on UDP traffic.
- Only the first 50 entries matching the selected statistic are listed. Therefore, if on the selected period of time, 100 pieces of information are identical, the GUI only displays the first 50.
- You might slow your appliance down if you edit the purge mechanism to include more lines or keep data longer than the default 30 days.

Monitoring DNS Queries and Answers

You can toggle the logs and monitor DNS physical servers queries and answers.

Monitoring DNS Queries

At any time you have the possibility to display all the DNS queries of an EfficientIP or a BIND DNS server. You can execute the command `querylog` from the page *All servers* and then display the whole list of logs on the page *Syslog* in the Administration module.

Querylog is a toggle command that provides an overview of all the DNS queries in IPv4 and IPv6. The logs structure is as follows:

- The requesting client IP address and port number, the query name, class and type.
- The recursion detailed, + or -. The Recursion Desired flag: + is set (the query was recursive), - is not set (the query was iterative).
- DNS options details if relevant: whether the query is signed (S), whether EDNS was used (E), whether TCP was used (T), whether DO - *DNSSEC OK* - was set (D), whether CD - *Checking Disabled* - was used (C).
- The IP address the information was sent to.

Keep in mind that all the logs can be displayed in the page *Syslog* in real time. They can slow this page down consistently as the `querylog` command can generate a substantial volume of data very quickly.

To toggle on the DNS querylog command

- Enabling the querylog**
 - In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
 - Tick the server of your choice.
 - In the menu, select **Edit > Command > Querylog**. The wizard **Toggle the querylog command** opens.
 - Click on to complete the operation. The report opens and closes. The page **All servers** is visible again and the server is marked **Enabled** in the column **Querylog**.

2. Displaying the DNS query and answer logs

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
- c. Under the menu, in the drop-down list **SOLIDserver**, make sure the hostname of the appliance managing the DNS server for which you toggled on the querylog is selected.
- d. In the drop-down list **Services**, select *named*.
- e. You can tick the box **Automatic refresh** if you want the page Syslog to refresh the log display every 10 seconds.
- f. Filter the list via the column **Time** with the current date, and the time if you want.
- g. In the column **Log**, all the query and answer logs of your server are displayed.

The first logs are *received control channel command 'querylog'* and *query logging is now on*; all the logs are listed below.

To toggle off the DNS querylog command

1. Disabling the querylog

- a. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
- b. Tick the server of your choice. It is marked *Enabled* in the column **Querylog**.
- c. In the menu, select  **Edit > Command > Querylog**. The wizard **Toggle the querylog command** opens.
- d. Click on **OK** to complete the operation. The report opens and closes. The page **All servers** is visible again and the server is marked **Disabled** in the column **Querylog**.

2. Making sure the querylog is off

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
- c. Under the menu, in the drop-down list **SOLIDserver**, make sure the hostname of the appliance managing the DNS server for which you toggled on the querylog is selected.
- d. In the drop-down list **Services**, select *named*.
- e. You can tick the box **Automatic refresh** if you want the page Syslog to refresh the log display every 10 seconds.
- f. Filter the list via the column **Time** with the current date, and the time if you want.
- g. In the column **Log**, two lines indicate the querylog is toggled off: *received control channel command 'querylog'* and *query logging is now off*.

Monitoring DNS Answers

At any time you have the possibility to display all the DNS answers of an EfficientIP or a BIND DNS server from the CLI.

Answerlog is a toggle command that provides an overview of all the DNS answers in IPv4 and IPv6. It uses the same information structure as *querylog*, each log contains: client information, recursion details, DNS options when relevant... For more details, refer to the [log structure](#) in the section [Monitoring DNS Queries](#).

When *answerlog* is on, it displays two lines: first the initial query, second the answer received by the client followed by the return code of the answer: *NOERROR*, *SERVFAIL*, *NXDOMAIN*...

To toggle on the DNS answerlog command

1. Connect to your appliance via a shell session.
2. Use the following command:

```
/usr/local/nessy2/bin/rndc answerlog
```

The logs look as follows:

```
# Line 1: the initial client query.
# Line 2: the answer sent to the client followed by the return code of the answer.
Nov 18 18:05:08 my-appliance-hostname named[1552]: client 127.0.0.1#31070 (www.google.com): view
vue3: query: www.google.com IN A +E (127.0.0.1)
Nov 18 18:05:08 my-appliance-hostname named[1552]: client 127.0.0.1#31070 (www.google.com): view
vue3: answer: www.google.com IN A +E (127.0.0.1) -> NOERROR www.google.com. 95 A 74.125.71.99
A 74.125.71.103 A 74.125.71.104 A 74.125.71.105 A 74.125.71.106 A 74.125.71.147
```

To toggle off the DNS answerlog command

1. Connect to your appliance via a shell session.
2. Use the following command:

```
/usr/local/nessy2/bin/rndc answerlog
```

Generating DNS Reports

EfficientIP provides dedicated DNS reports at server and zone level. The reports on inconsistencies or misconfiguration details might be empty if the server or zone configuration is correct.

Table 47.4. Available DNS reports

Page	Report
All servers	A records without an IPv4 address or alias
	AAAA records without an IPv6 address or alias
	CNAME records without an alias
	PTR records without an IPv4 address
	PTR records without an IPv6 address
	Route 53 Incompatibilities
	Zones NS and IP addresses
	Servers Configuration
	Hybrid DNS Engine Incompatibilities
	Server Peak Hour
	Servers Configuration Comparison
	Query Rate per Server
	Server Reply to Queries Charts

Page	Report
	Server Usage Charts
All views	View Statistics
All zones	Zones NS and IP addresses
	Zones Missing RRs
	Zone Statistics
	Zones Configuration Comparison

For more details regarding the reports and their generation, refer to the section [Managing Reports](#).

Part VIII. Global Policies

Global policies are options that allow to set specific behaviors within a module or between modules. They are available in the modules IPAM, DHCP, DNS and Device Manager.

- [Inheritance and Propagation](#) allows to can use meta-data, advanced properties and class parameters across the IPAM, DNS and Device Manager hierarchy. You can set, inherit, propagate or restrict an object's property from one level to the other.
 - [Managing Advanced Properties](#) allows to configure advanced properties that define interactions between and/or within the IPAM, DHCP and DNS modules.
-

Chapter 48. Inheritance and Propagation

Two mechanisms allow to inherit and propagate meta-data, advanced properties and class parameters within all modules. Both mechanisms are available by default in the addition/edition wizard of the resources that can be configured with it.

For each meta-data, advanced property and class parameter, you can configure the **Inheritance property** to *Set* or *Inherit* and/or the **Propagation property** to *Propagate* or *Restrict* via a dedicated icon and layer in the wizard.

Keep in mind that the default configuration of both properties does automatically *Propagate* the parameter value you specify on an object to all its child objects if their Inheritance property is *Inherit*. Like in the example below, at the space level, a DNS server is *Set*. All the objects the space contains down to its IP addresses *Inherit* this DNS server.

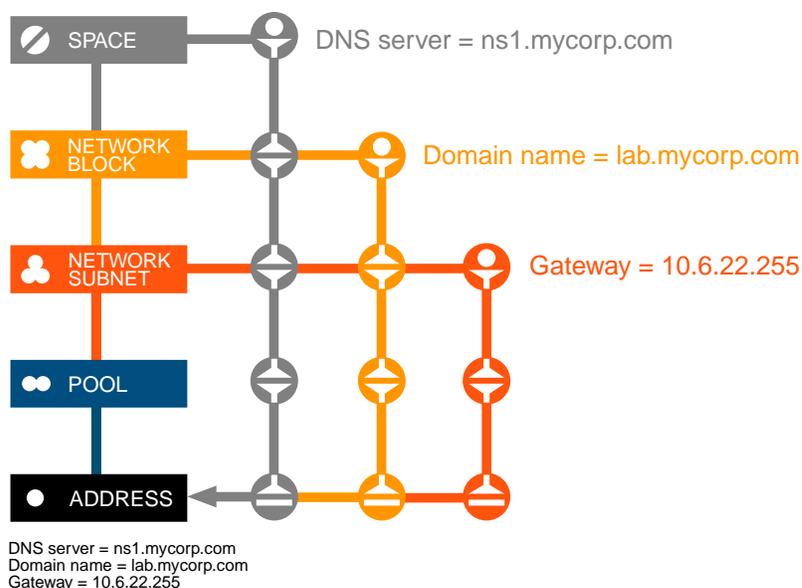


Figure 48.1. Example of the inheritance and propagation of some advanced properties in the IPAM

In the wizard, the dedicated configuration icon is located next to all the fields for which you can configure either the inheritance or the propagation property. It indicates the current configuration of the meta-data, advanced property or class parameter. Here below are the possible combinations.

Table 48.1. Inheritance and Propagation property possible configurations

Icon	Current Inheritance and Propagation configuration
☐	Inheritance: <i>Set</i> . Propagation: <i>Propagate</i> . The default value for both properties.
⊖	Inheritance: <i>Set</i> . Propagation: <i>Restrict</i> .
☐	Inheritance: <i>Inherit</i> . Propagation: <i>Propagate</i> .
⊖	Inheritance: <i>Inherit</i> . Propagation: <i>Restrict</i> .

You can use the internal hierarchy of a module to inherit and propagate a parameter with the values that suit your needs. In the example below, the advanced property DNS server has one value from space to block-type network and different one from subnet-type network down to IP address level.

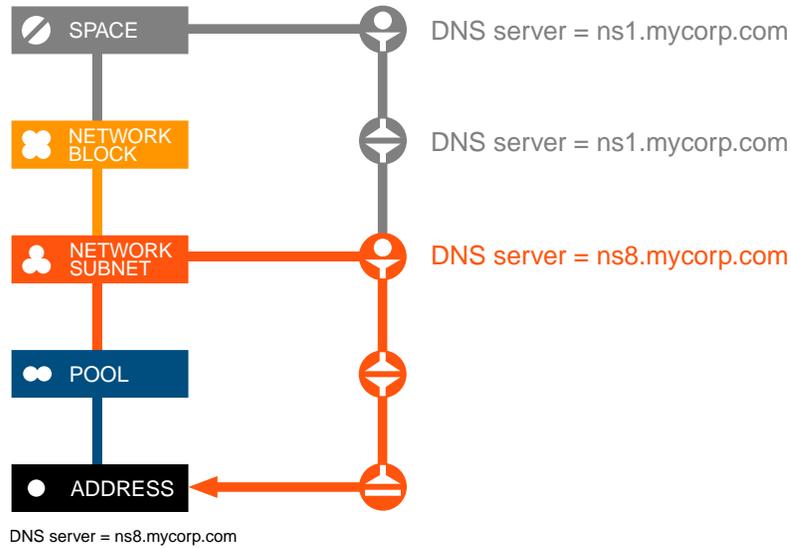


Figure 48.2. Example of an advanced property set and propagated at network level

Prerequisites

Before configuring the inheritance property and/or the propagation property on advanced properties, meta-data and class parameters, keep in mind that:

- **Advanced Properties** can be configured for the IPAM, DNS and/or DHCP. At all levels of these modules, you can define the inheritance and/or the propagation property of your advanced properties. For more details regarding advanced properties, refer to the chapter [Managing Advanced Properties](#).
- **Meta-data** is by default configured for class *global* of all resources within SOLIDserver. Note that a few resources of the modules IPAM, DNS and Device Manager provide meta-data for which you can define inheritance and propagation, without editing the class.

Table 48.2. Resources that have meta-data available by default

Module	Resource	Parameter
IPAM (IPv4 and IPv6)	Network (block-type)	Description
	Network (subnet-type)	
	Pool	
DNS	Zone (regular)	AD replication (Expert mode)
Device Manager	Device	Description

For more details regarding meta-data, refer to the chapter [Configuring Classes](#).

- **Class Parameters** belong to the customized classes you create from the page Class Studio. Once enabled, you can define the inheritance and propagation property of all the class objects that define your classes. Once displayed in the addition/editing wizard of the resource it applies to, the class objects are considered as class parameters. For more details regarding customized classes, refer to the chapter [Configuring Classes](#).

Limitations

- **Configuring the Inheritance and/or Propagation property does not trigger any creation.**

Configuring the inheritance and propagation properties on objects managing other objects propagates the parameter and its value to the objects it contains but does not trigger any creation behavior. So if you configure the advanced property *Create DHCP static* on a terminal network, the property value is propagated down to its assigned IP addresses but these addresses are not created in the DHCP.

To trigger the creation behavior on existing objects at lower level, you must tick them and in the menu select **Tools > Expert > Initialize rules**.

- **The Inheritance property cannot be configured at some levels:**
 - At the highest level of any module, you cannot switch the property to *Inherit*, it is forced to *Set*.
 - You cannot inherit a parameter on: IPAM spaces¹, DHCP servers, DNS servers, NetChange devices, Workflow requests, Device Manager devices, VLAN domains and groups of users.
- **The Propagation property cannot be configured at some levels:**
 - At the lowest level of any module, you cannot set the property to *Propagate*, it is forced to *Restrict*.
 - You cannot propagate parameters on: IP addresses (IPv4 and IPv6), DHCP groups (IPv4 and IPv6), DHCP ranges (IPv4 and IPv6), DHCP statics (IPv4 and IPv6), DNS zones, Device Manager interfaces and ports, NetChange ports and VLAN ranges. All these objects can however inherit parameters.
- **Some operations are impossible.** You cannot:
 - Configure the Inheritance property and Propagation property to *Inherit/Restrict* in template mode in the IPAM. An inherited parameter value must be propagated.
 - Delete a class parameter from a parent object if its Inheritance property is *Inherit*.
 - *Restrict* the propagation of a parameter that has already been propagated.
- **Reconciling class parameter must follow a specific order.**

If you plan on reconciling the meta-data, advanced properties and/or class parameters of all the objects of a module, you must start from the lowest level of the hierarchy up to the highest one.

Configuring the Inheritance of a Parameter Value

The inheritance property can be configured as follows:

- This property has two possible values: **Inherit** and **Set**.
- The default value of the property is *Inherit* if on the parent object the Propagation property is *Propagate*.
- This property is forced to *Set* if the parameter is not configured on the parent object.
- This property is forced to *Set* if the parameter is configured on the parent object but its Propagation property is *Restrict*.

¹If you have a space-based VLSM organization, you cannot inherit parameters on the top level space. You can inherit parameters from the first level down.

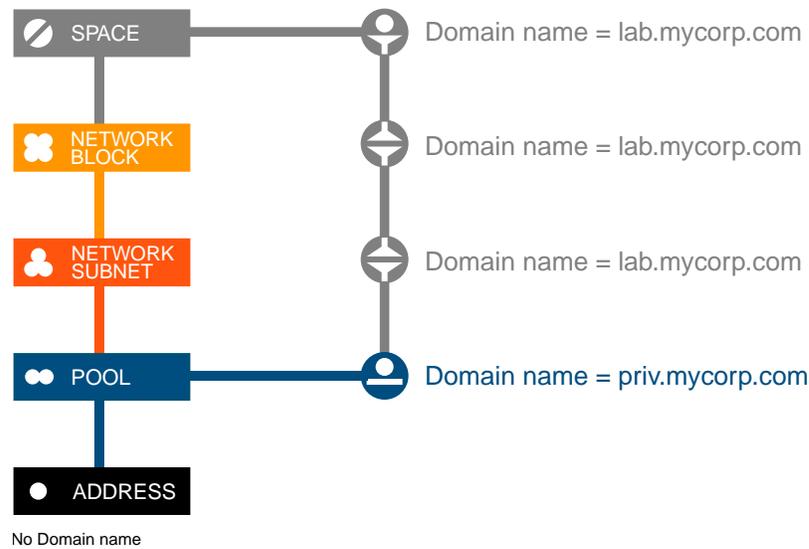


Figure 48.3. Example of an advanced property set and restricted at pool level

Once configured on a parent object, the value is used on all the objects it contains if their inheritance property is set to *Inherit*. In the example above, the value of the advanced property Domain name follows this logic: it is configured to *Set/Propagate* from the space down to the pool, every level in between is configured to *Inherit/Propagate*. At pool level the parameter is configured to *Set/Restrict* with a different value, so the IP addresses are not configured with the property.

In the following procedure the icon  is used to provide an example of configuration, it matches the default value of the inheritance and propagation properties. Depending on your configuration, it could be any of the icons detailed in the previous paragraph.

The inheritance property has to be configured directly in the addition/editing wizard of an object. In this section, *configuring* includes defining the inheritance property for the first time or editing its value.

To configure the inheritance of an advanced property or class parameter

1. Take into account the inheritance [Limitations](#).
2. Go to the module and page of your choice.
3. Add or edit a resource. The wizard opens.
4. If you or your administrator created classes, in the list **<object> class** select a class or *None*. Click on **NEXT**. The next page opens.
5. Right of the field of your choice, click on . The configuration layer opens.
6. In the drop-down list **Inheritance property**, click on *Set* or *Inherit* depending on your needs.

Keep in mind that you cannot configure the inheritance property on the highest level of a module.

7. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

The inheritance/propagation property icon details the object configuration on its properties page.

Note that deleting some objects of the IPAM and DNS hierarchy triggers the automatic inheritance:

- When you delete a non-terminal subnet-type network, an IPAM pool or a DNS view that contains other objects at lower levels, the child objects are automatically moved higher in the hierarchy.
- If the child objects inherited class parameters from the deleted container, for each class parameter:
 - The **Inheritance property** is forced to *Inherit* or *Set* to match the configuration of the deleted parent object. This way, the value and the source of the value remain the same.
 - The **Propagation property** remains the same.

Configuring the Propagation of a Parameter Value

The propagation property can be configured as follows:

- This property has two possible values: **Propagate** and **Restrict**.
- The default value of the property is *Propagate*. So the parameter value is propagated to the child objects:
 - If the parameter is not configured on the child objects.
 - If the parameter is not already configured with a different value.
 - If the parameter is not already configured with an Inheritance property *Inherit*.
- This property is forced to *Restrict* if it is impossible to propagate parameters on the child objects.

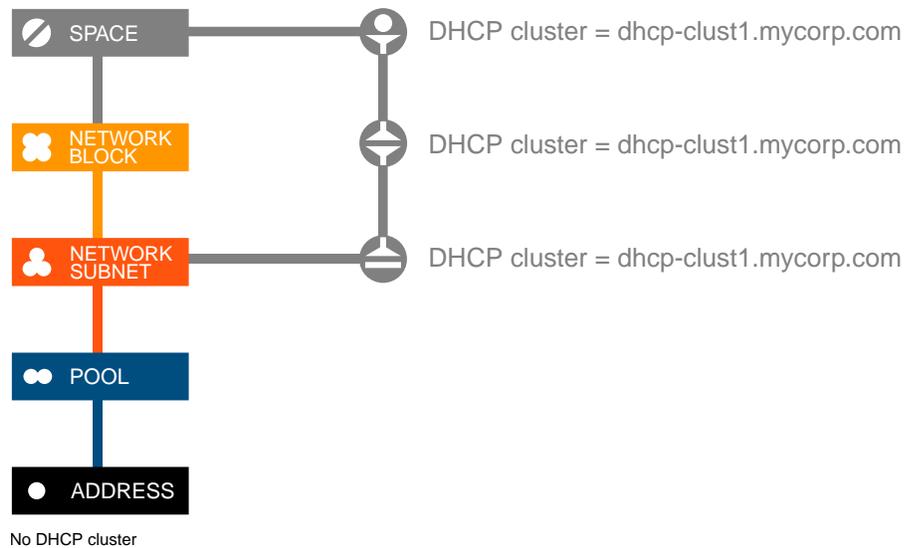


Figure 48.4. Example of an advanced property inherited and restricted at network level

Once configured on a parent object, the value is used on all the objects it contains if their propagation property is set to *Inherit*. In the example above, the value of the advanced property DHCP cluster follows this logic: it is configured to *Set/Propagate* from the space down to the subnet-type network, every level in between is configured to *Inherit/Propagate*. At network level the parameter is configured to *Inherit/Restrict*, so the pools and IP addresses are not configured with the property.

In the following procedure the icon  is used to provide an example of configuration, it matches the default value of the inheritance and propagation properties. Depending on your configuration, it could be any of the icons detailed in the previous paragraph.

The propagation property has to be configured directly in the addition/editing wizard of an object. In this section, *configuring* includes defining the propagation property for the first time or editing its value.

To configure the propagation of an advanced property or class parameter

1. Take into account the propagation [Limitations](#).
2. Go to the module and page of your choice.
3. Add or edit a resource. The wizard opens.
4. If you or your administrator created classes, in the list **<object> class** select a class or *None*. Click on **NEXT**. The next page opens.
5. At the right of the field of your choice, click on . The configuration layer opens.
6. In the drop-down list **Propagation property**, click on *Propagate* or *Restrict* depending on your needs.

Keep in mind that you cannot configure the propagation property on the lowest level of a module.

7. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

The inheritance/propagation property icon details the object configuration on its properties page.

Setting Class Parameters

You can set class parameters for several objects at once using the option *Set class parameters*. This option also allows to edit or define the inheritance and propagation configuration of the parameters of an object. This operation is quite advanced and should not be delegated lightly to users that do not belong the group *admin*.

If you only want to set advanced properties on several objects at once, we recommend that you use the option *Set advanced properties* because the wizard already contains all the available values and allows to choose them in dedicated drop-down lists rather than specifying them in fields. For more details, refer to the section [Setting Advanced Properties](#).

Using the option *Set class parameters* allows, as long as you are aware of the [Limitations](#), to:

- Set any parameter on an object: either a class parameter or meta-data.
- Overwrite any parameter configured or not.
- Overwrite the value of the Inheritance property and/or Propagation property of a parameter.
- Propagate a configured parameter to child objects if their Inheritance property is *Inherit*.
- Restrict a parameter at the level of your choice.

Note that, in the wizard, the drop-down lists *Inheritance property* and *Propagation property* are only displayed if they are relevant to the object you selected.

To set parameters inheritance and/or propagation property on several objects

1. Go to the module and page of your choice.
2. Tick the object(s) for which you want to set the inheritance and/or propagation property of a parameter.
3. In the menu, select **Tools > Expert > Set class parameters**. The wizard **Set class parameters** opens.
4. In the drop-down list **Parameter**, select the class parameter or meta-data of your choice. The page refreshes and the drop-down list **Value** appears.
5. In the drop-down list **Inheritance property**, configure the inheritance behavior:
 - a. Select **Set** to set the parameter the value or overwrite the current value on the object. This value is selected by default.
 - b. Select **Inherit** to inherit the advanced property from a parent object. Selecting this value hides the field *Value*.
6. If you selected *Set*, in the field **Value**, type in the parameter value.
7. In the drop-down list **Propagation property**, configure the propagation behavior:
 - a. Select **Propagate** to propagate the parameter value on the child objects. The propagation is only possible if the child objects Inheritance property is *Inherit*. This value is selected by default.
 - b. Select **Restrict** to only configure the parameter on the object(s) selected and prevent its propagation to the child objects.
8. Click on **ADD** to move your parameter configuration to the **Parameters list**.
 - To update an entry, select a configuration in the list, it is loaded again in the fields. Edit the needed data and click on **UPDATE**.
 - To delete an entry, select a configuration and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
9. Repeat the steps 5 to 8 for as many parameters as you need.
10. Click on **OK** to complete the operation. The report opens and closes. Any parameter value and configuration previously set for the selected object(s) is overwritten.

Reconciling Class Parameters

The option *Reconcile class parameters* allows to adjust the inheritance property of your meta-data, advanced properties and class parameters on several objects at once.

For instance, after a migration of your database, from a version prior to 6.0.0, with parameters configured that are all *Set* with the same value at all levels, you can use the option to make sure they respect the internal module hierarchy and make the lower levels *Inherit* the top level parameter value. If you configured or edited classes, your objects are configured at each level but the inheritance between each level is not implemented yet. Running this option allows comparing all the properties and parameters set on the parent, for the selected child objects. If they are configured on both levels and their values match, the inheritance property of the child objects is forced to *Inherit* to set up the inheritance.

The option allows to force the *Inheritance property* of a parameter based on the value of the parameter in the parent object. Therefore:

- The parameter value is forced to *Set* if it is not configured on the parent object.
- The parameter value is forced to *Set* if it is configured on the parent object with the Propagation property set to *Restrict*.
- The parameter value is forced to *Set* if it has a different value on the parent object.
- The parameter value is forced to *Inherit* if on the parent object: the parameter is configured, has the same value and has the Propagation property set to *Propagate*.

Keep in mind that the option is only available on pages managing objects that can inherit data. Some class parameters might not be reconciled, for more details refer to the section [Limitations](#).

To reconcile a parameter inheritance property

1. Go to the module and page of your choice.

If you plan on reconciling meta-data, advanced properties and class parameters on all the objects of a module, you must start from the lowest level and execute the option on all levels up to the highest one.

2. Tick the object(s) for which you want to reconcile the class parameters inheritance property with their parent object.
3. In the menu, select  **Tools > Expert > Reconcile the class parameters**. The wizard **Reconcile the class parameters** opens.
4. Click on  to complete the operation. The report opens and works for a while. The report opens.

Chapter 49. Managing Advanced Properties

Advanced properties are specific class parameters that allow administrators to configure replication behaviors between the modules IPAM, DHCP and DNS.

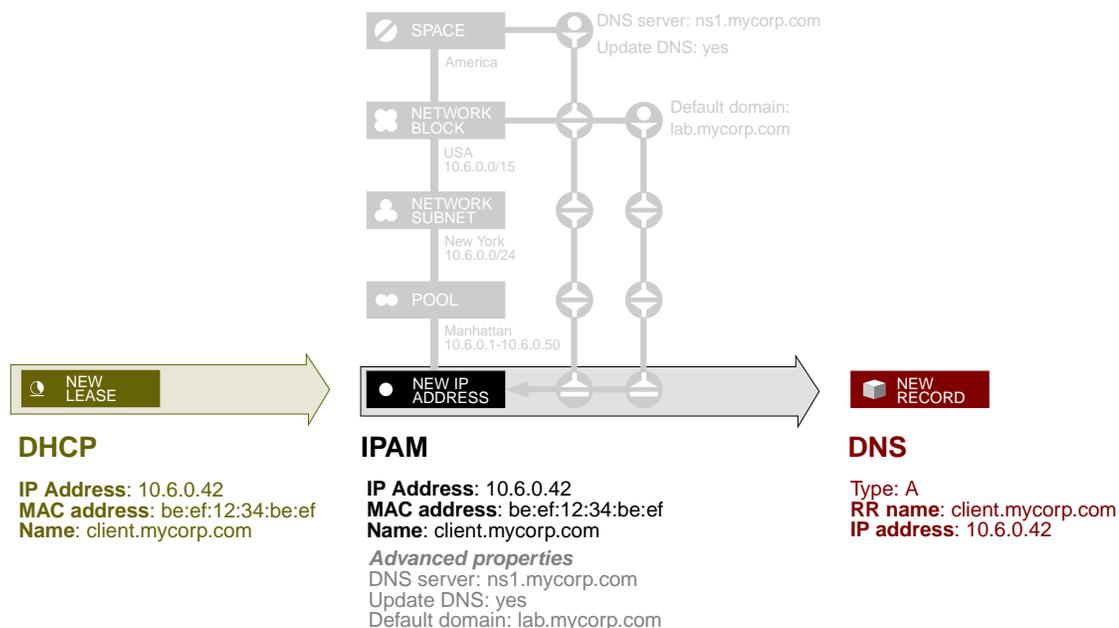


Figure 49.1. Replication from the DHCP to the DNS via the IPAM

These properties can edit the interaction between these three modules and/or provide extra options within a module:

- **From the IPAM module**, you can update the DNS and/or DHCP when creating networks or adding IP addresses, within this module you can also define the gateway creation behavior or decide to create pools when you create terminal networks.

The module also provides extra advanced options that are detailed in separate chapters:

- The IPv4 to IPv6 transition properties are detailed in the chapter [Setting Up a Transition From IPv4 to IPv6](#).
- The IPAM/Device Manager interaction properties are detailed in the chapter [Managing the Interaction with the IPAM](#).
- The IPAM/VLAN interaction properties are detailed in the chapter [Managing the IPAM/VLAN Interaction](#).
- **From the DHCP module**, you can update the IPAM and DNS when allocating leases.
- **From the DNS module**, you can update the IPAM when updating the records database of a smart architecture.

Keep in mind that once configured at high level, the value of advanced properties is inherited by all the objects at lower level, unless you configure them otherwise. Which is why, you can configure

on a space properties that apply to IP addresses and you can configure on a server properties that apply to leases or records.

The advanced properties configuration involves two operations:

1. Selecting at each level the advanced properties fields to display or enable, by default, in the addition and edition¹ wizards.
2. In the addition or edition wizard, configuring the property. The configuration of the properties implies that the end user has sufficient rights and resources.

All the advanced properties are represented in the appendix [Advanced Properties](#).

You can edit the advanced properties inheritance and propagation of several objects at a time. For more details, refer to the section [Setting Advanced Properties](#).

Note that users belonging to a group that do have the right *Advanced properties Customize: wizard* can view but cannot edit the advanced properties parameters applied to an object. They appear in gray in the addition and edition wizards.

Prerequisites

To use advanced properties you must select all the modules in the wizard *Internal module setup*. It is accessible:

- From the page **Main dashboard**, in the gadget *SOLIDserver configuration checklist*, next to *Internal module setup* click on **Configuration**.
- From the page **Admin Home**, in the section *Expert*, click on **Internal module setup**.

For more details, refer to the section [Defining the Internal Module Setup](#).

Once the internal module setup is complete, administrators must select the properties they want to display in the addition/edition wizard of the resources. For more details, refer to the section [Selecting the Advanced Properties Displayed by Default](#).

Browsing Advanced Properties

There is no page dedicated to advanced properties, but the objects that can be configured with them contain the panel **Advanced properties**.

That panel displays the advanced properties configuration of the object. After the property value, you may find between brackets the level the property was inherited from.

¹Note that the edition wizard of DNS and DHCP servers managed via a smart architecture does not display the advanced properties fields. They can only be edited from the smart architecture itself.

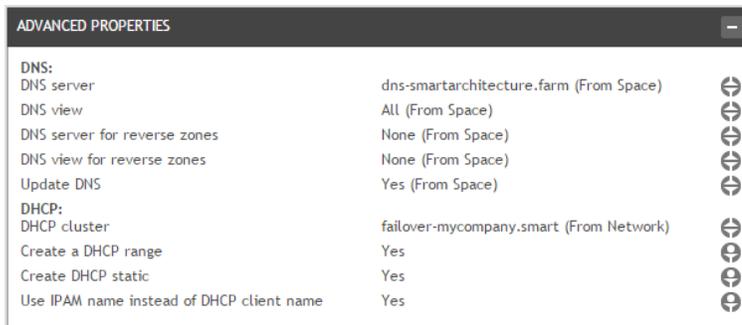


Figure 49.2. The panel Advanced properties

To display the panel Advanced properties

1. Go to the properties page of the object of your choice using .
2. Open the panel **Advanced properties**, using . The panel displays the advanced properties configuration of the object.

Selecting the Advanced Properties Displayed by Default

End users with limited rights can only configure advanced properties on IPAM, DNS and/or DHCP resources if their administrator chose to display them.

That display configuration can only be set via the wizard *Advanced properties customization*. Once the wizard is configured:

- The selected properties are displayed and/or enabled by default in the resource addition/edition wizard.
- At the bottom of the addition/edition wizard, the drop-down list **Advanced properties** possible values are:
 - **Default:** all the properties that were ticked/enabled by the administrator in the wizard *Advanced properties customization* are displayed. This is the drop-down list default value.
 - **All:** all the properties available for display in the wizard *Advanced properties customization* are loaded. That option value excludes the properties that must be enabled in the customization wizard. Only administrators or users with sufficient rights can select *All* in the list.

To select the advanced properties to display in the addition/edition wizard

1. Depending on your needs, in the sidebar:
 - a. Go to  **IPAM** > **Networks**, **Pools** or **Addresses**. The page opens.
 - b. Go to  **DHCP** > **Servers**, **Scopes** or **Ranges**. The page opens.
 - c. Go to  **DNS** > **Servers**, **Views**, or **Zones**. The page opens.
2. In the menu, select  > **Extra options** > **Wizard customization**. The wizard **Advanced properties customization** opens.
3. Tick the boxes of the properties you want to display or enable in the addition/edition wizard.

Each object has its own properties, refer to the module and object that suits your needs.

- Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

In the addition/edition wizard, all the ticked and enabled properties are now part of the *Default* display.

Configuring IPAM Advanced Properties

To configure advanced properties:

- Your administrator must have configured the internal module setup as detailed in the [Prerequisites](#) and chose the fields available for configuration as detailed in the section [Selecting the Advanced Properties Displayed by Default](#).
- You must configure the properties in the addition/edition wizard of the object.

You can configure **IPAM spaces, networks, pools and IP addresses managing IPv4 and IPv6 addresses** with IPAM properties, DNS properties and/or DHCP properties.

Note that within the IPAM, two options use the same principles than the advanced properties but are not considered as such because they only apply to IPAM objects. For more details refer to the chapters [Setting Up a Transition From IPv4 to IPv6](#), [Managing the Interaction with the IPAM](#), and [Managing the IPAM/VLAN Interaction](#).

Space Advanced Properties

At space level, the wizard *Advanced properties customization* only contains advanced properties that can be displayed. All properties are enabled.

The table below details the entries of the customization wizard and the corresponding field in the addition/edition wizard.

Table 49.1. Space advanced properties fields

In the wizard Advanced properties customization	In the addition/edition wizard
Display the IPv4 to IPv6 transition fields	Activate the IPv4 to IPv6 transition
Select the DNS server where the IP addresses should be updated	DNS server
Select the DNS view where the IP addresses should be updated	DNS view
Select the DNS domain where the IP addresses should be updated	Default Domain
Select a restricted list of the allowed domains	Domains list / Selected Domains list
Create a DNS reverse zone	<i>Tick the box to display more fields</i>
Select the DNS server where the reverse zone of a network should be created	DNS server for reverse zones
Select the DNS view where the reverse zone of a network should be created	DNS view for reverse zones
Select the DHCP failover cluster where the configuration should be applied	DHCP cluster ^a
Display the box "Update DNS"	Update DNS
Display the box "Create DHCP static"	Create DHCP Static

^aSetting a DHCP cluster at space level is only propagated to its IPV4 objects.

Once the advanced properties are displayed, they can be configured when adding or editing spaces. Each one can be inherited at lower levels to automate the replication from the IPAM.

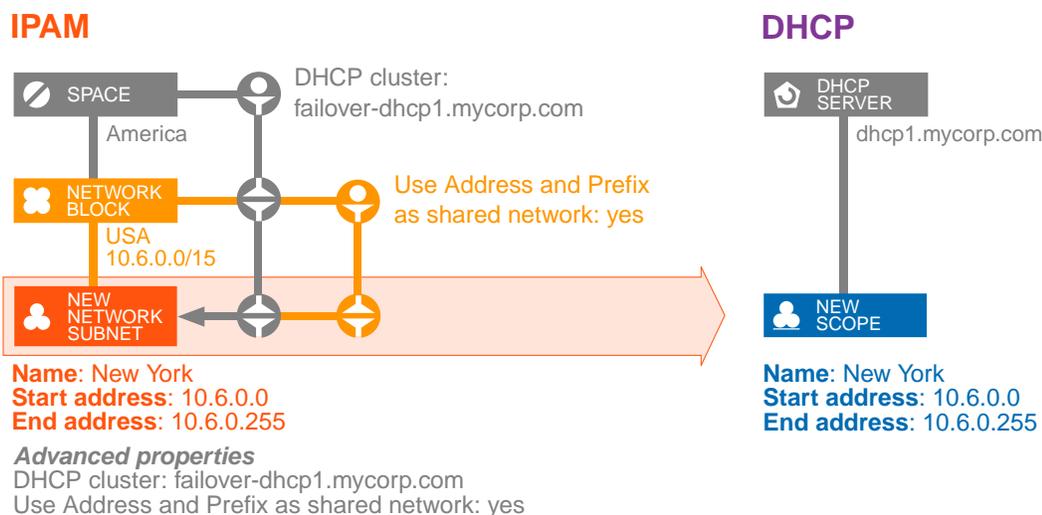


Figure 49.3. IPAM to DHCP replication at network level inherited from the space

To configure advanced properties at space level

1. In the sidebar, go to **IPAM > Spaces**. The page **All spaces** opens.
2. Add or edit a space. The wizard opens. For more details refer to the chapter [Managing Spaces](#).
3. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on **⚙** to define its *Propagation property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.2. Space advanced properties configuration fields

Field	Default value	Description
Activate the IPv4 to IPv6 transition	Unticked	Tick this box to configure the IPv4 to IPv6 transition. For more details, refer to the chapter Setting Up a Transition From IPv4 to IPv6 .
DNS server	None	Select the DNS server managing the zone you want to update when you assign IP addresses. Once you selected a server, the fields <i>DNS view</i> , <i>Domains list</i> , <i>Selected domains list</i> and <i>Default domain</i> appear if they were ticked in the customization wizard. If you select <i>All</i> , all the views and zones are available in the related fields. You must tick the box Update DNS to take this parameter into account.
DNS view	All	Select the DNS view managing the zone you want to update when you assign IP addresses. You must tick the box Update DNS to take this parameter into account.
Domains list	/	Select the DNS domain, i.e. zone, you want to update when you assign IP addresses. You might have the possibility to add several zones, use + to move them to the field <i>Selected domains</i>

Field	Default value	Description
		<i>list</i> . You must tick the box Update DNS to take this parameter into account.
Default domain	None	Select a default domain. The listed zones are all part of the list <i>Selected domains</i> . You must tick the box Update DNS to take this parameter into account.
DNS server for reverse zones	All	Select the DNS server where reverse zones are created upon addition/edition of terminal networks. Keep in mind that deleting a network configured with this option also deletes the corresponding reverse zone from the DNS, if it only contains the default records (SOA and NS). You must tick the box Update DNS to take this parameter into account.
DNS view for reverse zones	All	Select the DNS view where the reverse zones are created upon addition/edition of terminal networks. Keep in mind that deleting a network configured with this option also deletes the corresponding reverse zone from the DNS, if it only contains the default records (SOA and NS). You must tick the box Update DNS to take this parameter into account.
Update DNS	Unticked	Tick this box to update the zone configured as <i>default domain</i> with A, AAAA, PTR and/or CNAME records when you assign IP addresses in the relevant terminal network ^a . Keep in mind that these records, created from the IPAM, are deleted when you untick the box.
DHCP cluster	None	Select a DHCP failover cluster to create a scope matching every terminal network you create (name and addresses), on the DHCP servers of the selected cluster. The terminal network you create contains the panel <i>DHCP options</i> displaying the options configuration of the scope. In IPv4, the DHCP servers are automatically set with the option routers, its value is the network gateway address.
Create DHCP static	Unticked	Tick this box to reserve a DHCP static for every assigned IP address. The static creation is only possible if you selected a <i>DHCP cluster</i> .
Use IPAM name instead of DHCP client name	Unticked	Tick this box to name the static created in the DHCP like the IP address you assign in the IPAM. This static is automatically configured with the option <i>host-name</i> set to the value of the IP address <i>Shortname</i> , and if you selected a <i>Default domain</i> , the static is also configured with the option <i>domain-name</i> set to the value of the default domain you chose. If you do not tick it, the IP address you create does not have a name and it is named after the next DHCP client that connects to the network.

^aNote that the gateway address no longer creates an A or AAAA record in the DNS.

- Click on to complete the operation. The report opens and closes. The page **All spaces** is visible again.

Network Advanced Properties

At network level, the wizard *Advanced properties customization* contains advanced properties that can appear in the subnet-type and/or block-type network wizards.

Enabling the advanced properties can allow to automatically create a DNS zone when you add a network.

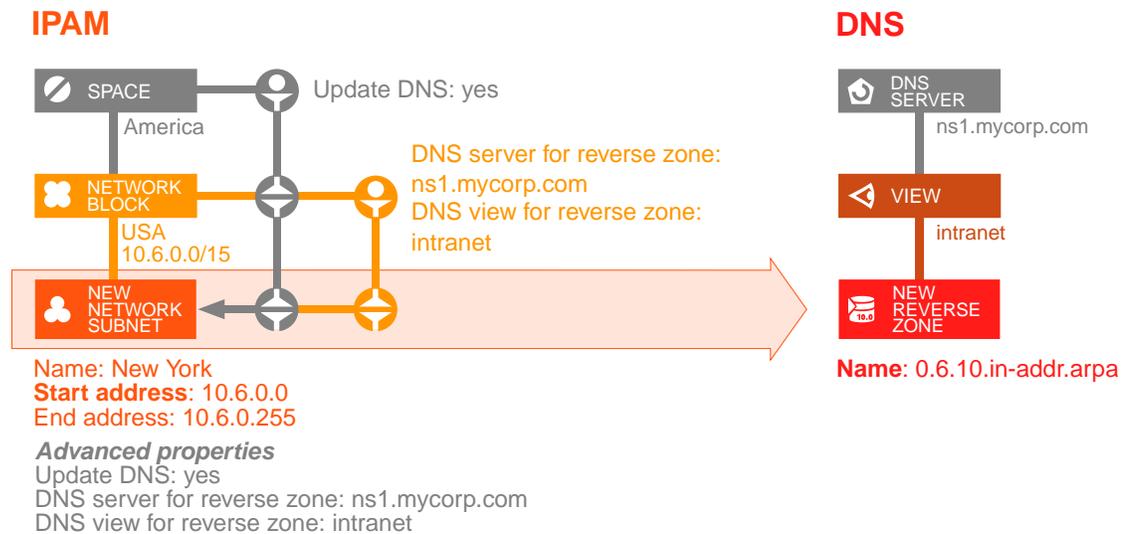


Figure 49.4. IPAM to the DNS replication at network level

The existing networks are not impacted, to configure them like their container refer to the section [Setting Advanced Properties](#).

The table below details the entries of the customization wizard and the corresponding field in the addition/editing wizard.

Table 49.3. Network advanced properties fields

In the wizard Advanced properties customization	In the addition/editing wizard
IPAM properties	
Gateway offset	Gateway
Display the field "Gateway"	Gateway
Display the IPv4 to IPv6 transition fields	Activate the IPv4 to IPv6 transition
DNS properties	
Select the DNS server where the IP addresses should be updated	DNS server
Select the DNS view where the IP addresses should be updated	DNS view
Select the DNS domain where the IP addresses should be updated	Default Domain
Select a restricted list of the allowed domains	Domains list / Selected Domains list
Create a DNS reverse zone	<i>Tick the box to display more fields</i>
Select the DNS server where the reverse zone of a network should be created	DNS server for reverse zones
Select the DNS view where the reverse zone of a network should be created	DNS view for reverse zones
Display the box "Update DNS"	Update DNS
DHCP properties	
Select the DHCP failover cluster where the configuration should be applied	DHCP cluster ^a
Ask if this network must be configured as a DHCP shared network ^b	Use Address and Prefix as shared network
Display the box "Create DHCP static"	Create DHCP Static
Pools creation	
Ask the number of pools to create	Number of pools / Size / Type / Name
IPAM/VLAN interaction	

In the wizard Advanced properties customization	In the addition/edition wizard
Display the VLAN association fields	VLAN domain
	VLAN range
	VLAN ID

^aIn IPv6 this property cannot be inherited from the server. You must set it to propagate it down to the IP addresses.

^bTick the box *Select the DHCP failover cluster where the configuration should be applied* to display this field.

Note that from the customization wizard, you can enable or disable two properties:

Gateway offset

This field allows to automatically calculate and create a gateway when you add terminal networks:

- You can specify positive and negative values in the field. A positive value calculates the gateway from the start address of the network; a negative value calculates it from the end address of the network.

By default, the offset value is *-1*: the penultimate address of the network is used as gateway.

- You cannot specify *0* in the field. To disable the automatic gateway creation and create terminal networks without a gateway, you must leave the field empty. Once the gateway creation is disabled, the field *Gateway* is not displayed in the addition/edition wizard even if the box *Display the field "Gateway"* is ticked.

Ask the number of pools to create

This box allows to create pools during the creation of a terminal network. Note that:

- This option does not apply when you edit a terminal networks.
- This option can only be used if pool classes are enabled.
- The field *Number of pools* cannot be displayed in the addition wizard if you do not tick this option.

As the properties are inherited, the advanced properties configuration at network level influence the IP address behaviors. That way, the automatic creation of the *Gateway* address, configures the DHCP option *Routers* on the scope automatically created when you add a terminal network.

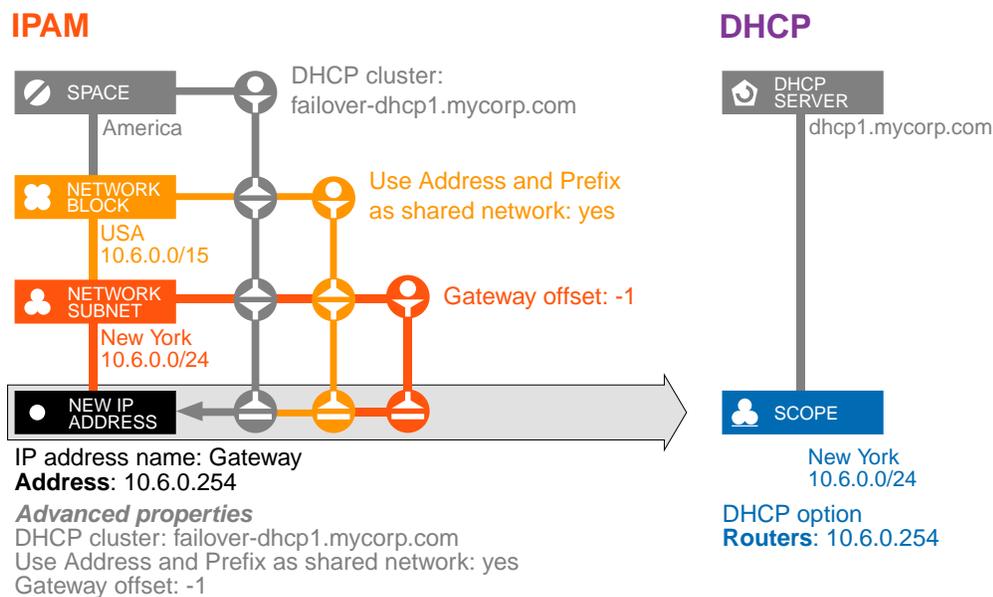


Figure 49.5. Replication from the IPAM to the DHCP at IP address level

All the properties detailed in the procedure below are based on a configuration that was not inherited. Most fields are already filled with the value in the container of the resource you create. For more details, refer to the chapter [Inheritance and Propagation](#).

To configure advanced properties at network level

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit a network. The wizard opens. For more details refer to the chapter [Managing Networks](#).
4. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on **⚙** to define its *Inheritance property* and/or *Propagation property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.4. Network advanced properties configuration fields

Field	Default value	Description
For all networks		
Activate the IPv4 to IPv6 transition	Unticked	Tick this box to configure the IPv4 to IPv6 transition. For more details, refer to the chapter Setting Up a Transition From IPv4 to IPv6 .
DNS server	None	Select the DNS server managing the zone you want to update when you assign IP addresses. Once you selected a server, the fields <i>DNS view</i> , <i>Domains list</i> , <i>Selected domains list</i> and <i>Default</i>

Managing Advanced Properties

Field	Default value	Description
		<i>domain</i> appear if they were ticked in the customization wizard. If you select <i>All</i> , all the views and zones are available in the related fields. You must tick the box Update DNS to take this parameter into account.
DNS view	All	Select the DNS view managing the zone you want to update when you assign IP addresses. You must tick the box Update DNS to take this parameter into account.
Domains list	/	Select the DNS domain, i.e. zone, you want to update when you assign IP addresses. You might have the possibility to add several zones, use + to move them to the field <i>Selected domains list</i> . You must tick the box Update DNS to take this parameter into account.
Default domain	None	Select a default domain. The listed zones are all part of the list <i>Selected domains</i> . You must tick the box Update DNS to take this parameter into account.
DNS server for reverse zones	All	Select the DNS server where reverse zones are created upon addition/edition of terminal networks. Keep in mind that deleting a network configured with this option also deletes the corresponding reverse zone from the DNS, if it only contains the default records (SOA and NS). You must tick the box Update DNS to take this parameter into account.
DNS view for reverse zones	All	Select the DNS view where the reverse zones are created upon addition/edition of terminal networks. Keep in mind that deleting a network configured with this option also deletes the corresponding reverse zone from the DNS, if it only contains the default records (SOA and NS). You must tick the box Update DNS to take this parameter into account.
Update DNS	Unticked	Tick this box to update the zone configured as <i>default domain</i> with A, AAAA, PTR and/or CNAME records when you assign IP addresses in the relevant terminal network ^a . Keep in mind that these records, created from the IPAM, are deleted when you untick the box.
DHCP cluster	None	Select a DHCP failover cluster to create a scope matching every terminal network you create (name and addresses), on the DHCP servers of the selected cluster. The terminal network you create contains the panel <i>DHCP options</i> displaying the options configuration of the scope. In IPv4, the DHCP servers are automatically set with the option routers, its value is the network gateway address.
Create DHCP static	Unticked	Tick this box to reserve a DHCP static for every assigned IP address. The static creation is only possible if you selected a <i>DHCP cluster</i> .
Use IPAM name instead of DHCP client name	Unticked	Tick this box to name the static created in the DHCP like the IP address you assign in the IPAM. This static is automatically configured with the option <i>host-name</i> set to the value of the IP address <i>Shortname</i> , and if you selected a <i>Default domain</i> , the static is also configured with the option <i>domain-name</i> set to the value of the default domain you chose. If you do not tick it, the IP address you create does not have a name and it is named after the next DHCP client that connects to the network.
Only for subnet-type networks		
Gateway	<i>Matches the offset</i>	Specify a different gateway if need be. Its value is calculated based on the <i>Gateway offset</i> configuration. If the gateway offset calculation is disabled, the field is not displayed.

Field	Default value	Description
Use Address and Prefix as shared network	Ticked	Tick this box to use the <i>Address</i> and <i>Prefix</i> of the IPv4 network you are adding as a shared network. This action creates a new scope in the DHCP servers of the selected cluster. This box is only visible when the <i>DHCP cluster</i> is <i>Set</i> and a value is specified.
Number of pools	0	Select between <i>1</i> and <i>5</i> pool(s) to create within the terminal network you are creating. For each pool, you need to specify a <i>Size</i> and <i>Type</i> . If you select <i>0</i> , no pool is created.
Size	None	Type in the number of IP addresses of each pool. This option is only available in IPv4.
Type	Select a value	Select a class for the pool. The pool is named after the class you select.
Name	None	If you selected <i>Other</i> , type in a name for the pool.
VLAN domain	None	Select a domain to associate a terminal network with a VLAN. For more details, refer to the chapter Managing the IPAM/VLAN Interaction .

^aNote that the gateway address no longer creates an A or AAAA record in the DNS.

- Click on to complete the operation. The report opens and closes. The page **All networks** is visible again.

Pool Advanced Properties

At pool level, the wizard *Advanced properties customization* only contains advanced properties that can be displayed. All properties are enabled.

Enabling the advanced properties allows to automatically create a DHCP range when you add a pool.

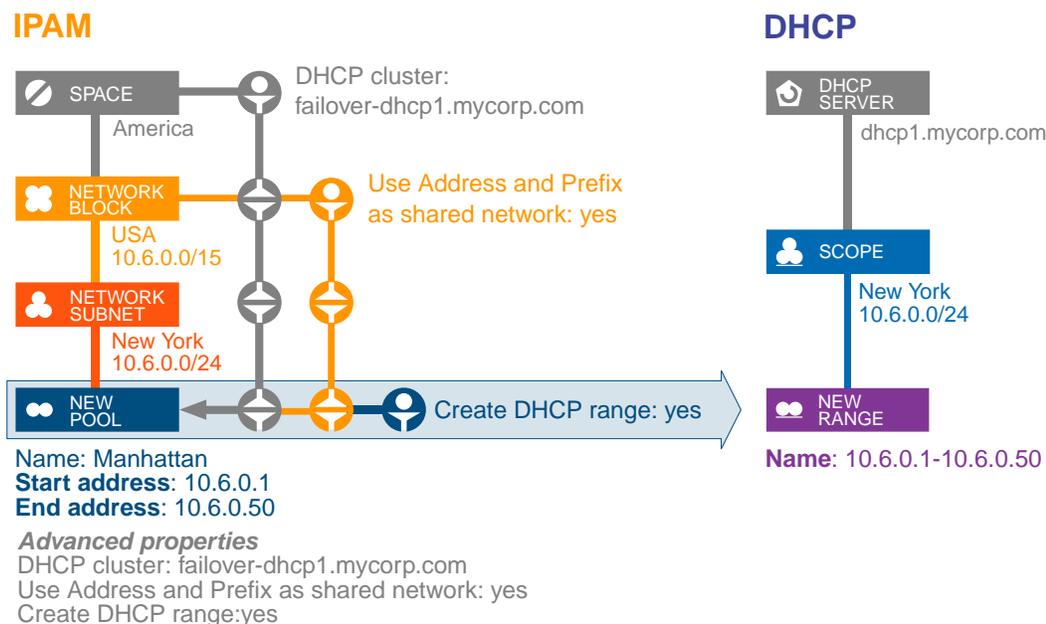


Figure 49.6. Replication from the IPAM to the DNS at pool level

The table below details the entries of the customization wizard and the corresponding field in the addition/edition wizard.

Table 49.5. Pool advanced properties fields

In the wizard Advanced properties customization	In the addition/edition wizard
Display the field "Create a DHCP range"	Create a DHCP range
Display the field "Create DHCP static"	Create DHCP static

All the properties detailed in the procedure below are based on a configuration that was not inherited. Most fields are already filled with the value in the container of the resource you create. For more details, refer to the chapter [Inheritance and Propagation](#).

To configure advanced properties at pool level

1. In the sidebar, go to **IPAM > Pools**. The page **All pools** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit a pool. The corresponding wizard opens. For more details, refer to the chapter [Managing Pools](#).
4. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on  to define its *Inheritance property* and/or *Propagation property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.6. Pool advanced properties configuration fields

Field	Default value	Description
Create a DHCP range	Unticked	Tick this box to create a range in the DHCP matching the start and end address of the pool you are creating. The range is created in the scope matching the terminal network the pool belongs to.
Create DHCP static	Unticked	Tick this box to reserve a DHCP static for every assigned IP address in the pool. The static creation is only possible if you selected a DHCP cluster.

5. Click on **OK** to complete the operation. The report opens and closes. The page **All pools** is visible again.

IP Address Advanced Properties

At IP address level, the wizard *Advanced properties customization* only contains advanced properties that can be displayed. All properties are enabled.

The table below details the entries of the customization wizard and the corresponding field in the addition/edition wizard.

Table 49.7. IP address advanced properties fields

In the wizard Advanced properties customization	In the addition/edition wizard
IPv4 addresses transition to IPv6	Corresponding IPv6 address

In the wizard Advanced properties customization	In the addition/edition wizard
Enable the automatic construction of the IP address hostname: short-name.domain	Shortname / Domain
Make the Domain selection mandatory for the hostname construction of the IP address and its aliases ^a	Domain (required)
Select the DNS server where the IP addresses should be updated	DNS server
Select the DNS view where the IP addresses should be updated	DNS view
Select the DNS domain where the IP addresses should be updated	Domain
Display the box "Update DNS"	Update DNS
Display the box "Create DHCP static"	Create DHCP Static
IPAM / Device Manager interaction^b	
Enable to create devices from the IPAM	Create a device
Enable to link IP address with existing devices	Interface name
Enable to edit the devices topology from the IPAM	Link with device / Link with port

^aThat box is only displayed if the box *Enable the automatic construction* is ticked.

^bFor more details, refer to the chapter [Managing the Interaction with the IPAM](#).

Enabling the advanced properties allows to automatically create a DHCP static or a DNS record when you add an IP address.

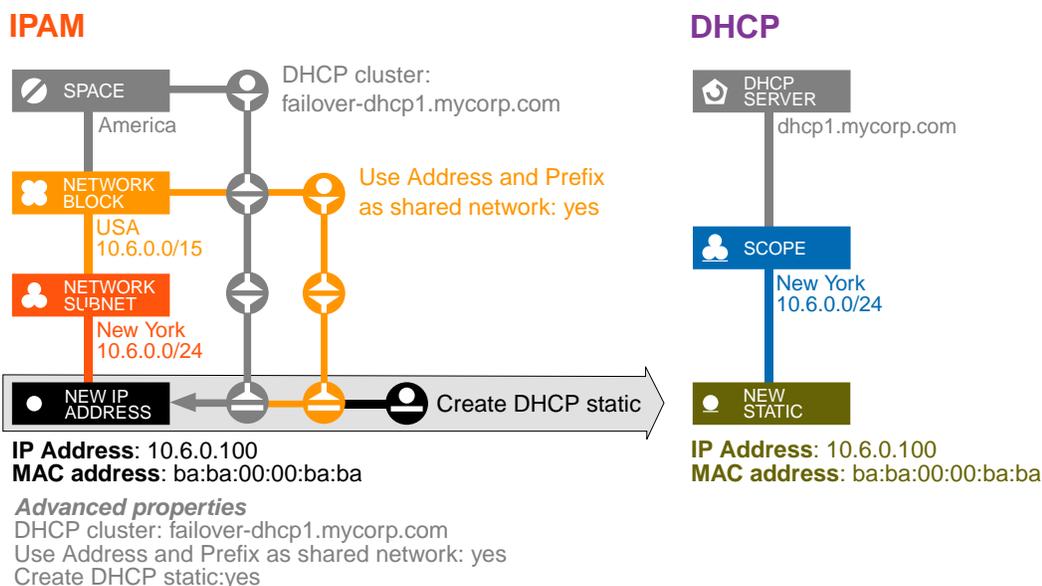


Figure 49.7. Replication from the IPAM to the DHCP at IP address level

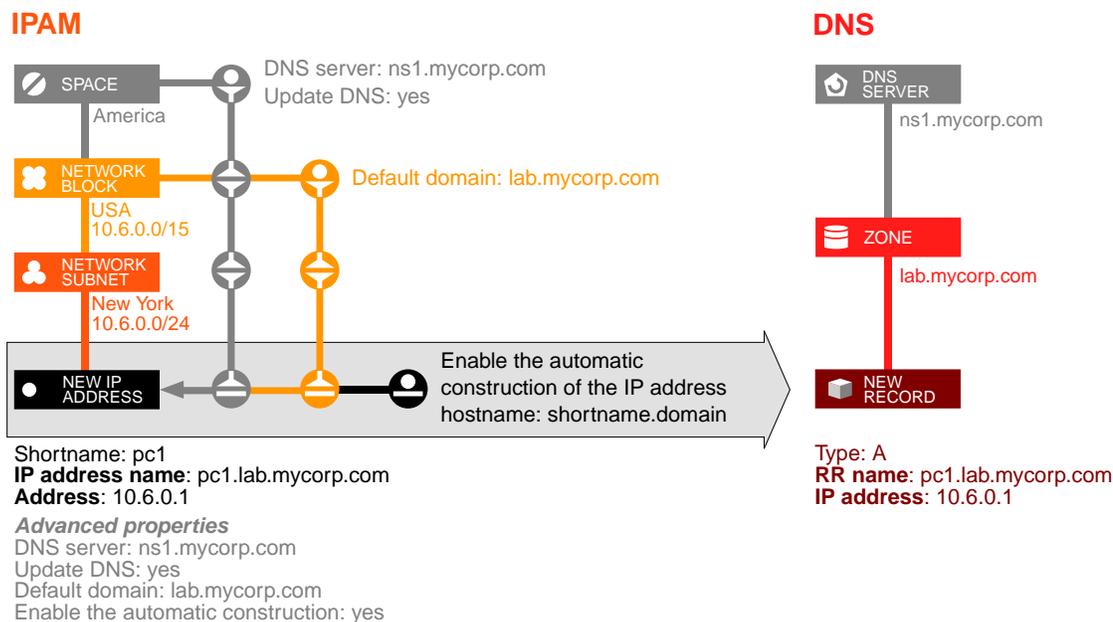


Figure 49.8. Replication from the IPAM to the DNS at IP address level

All the properties detailed in the procedure below are based on a configuration that was not inherited. Most fields are already filled with the value in the container of the resource you create. For more details, refer to the chapter [Inheritance and Propagation](#).

To configure advanced properties at IP address level

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit an IP address. For more details, refer to the chapter [Managing IP Addresses](#). The corresponding wizard opens.
4. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on **⚙** to define its *Inheritance property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.8. IP address advanced properties configuration fields

Field	Default value	Description
Corresponding IPV6 address	IP address	The field is read-only and matches the transition configuration. For more details, refer to the chapter Setting Up a Transition From IPv4 to IPv6 .
DNS server	None	Select the DNS server managing the zone you want to update when you assign IP addresses. Once you selected a server, the fields <i>DNS view</i> , <i>Domains list</i> , <i>Selected domains list</i> and <i>Default domain</i> appear if they were ticked in the customization wizard.

Field	Default value	Description
		If you select <i>All</i> , all the views and zones are available in the related fields. You must tick the box Update DNS to take this parameter into account.
DNS view	All	Select the DNS view managing the zone you want to update when you assign IP addresses. You must tick the box Update DNS to take this parameter into account.
Shortname	/	Type in a name for the IP address. The complete <i>IP address name</i> is constructed as follows: <i><shortname>.<domain_name></i> . Selecting this option creates an <i>A</i> record for the address in the DNS zone specified in the field <i>Domain</i> .
Domain	None	Select a domain to complete the <i>IP address name</i> . Selecting this option creates <i>A</i> , <i>AAAA</i> and/or <i>CNAME</i> records for the IP address and its aliases in the DNS zone specified in the field <i>Domain</i> . This field is mandatory if the box <i>Make the Domain selection mandatory in the hostname of the IP address and its aliases</i> is ticked in the customization wizard.
Update DNS	Unticked	Tick this box to update the zone configured as <i>default domain</i> with <i>A</i> , <i>AAAA</i> , <i>PTR</i> and/or <i>CNAME</i> records when you assign IP addresses in the relevant terminal network ^a . Keep in mind that these records, created from the IPAM, are deleted when you untick the box.
Create DHCP static	Unticked	Tick this box to reserve a DHCP static for every assigned IP address. The static creation is only possible if you selected a <i>DHCP cluster</i> .
Use IPAM name instead of DHCP client name	Unticked	Tick this box to name the static created in the DHCP like the IP address you assign in the IPAM. This static is automatically configured with the option <i>host-name</i> set to the value of the IP address <i>Shortname</i> , and if you selected a <i>Default domain</i> , the static is also configured with the option <i>domain-name</i> set to the value of the default domain you chose. If you do not tick it, the IP address you create does not have a name and it is named after the next DHCP client that connects to the network.

^aNote that the gateway address no longer creates an *A* or *AAAA* record in the DNS.

- Click on to complete the operation. The report opens and closes. The page **All addresses** is visible again.

Configuring DHCP Advanced Properties

To configure advanced properties:

- Your administrator must have configured the internal module setup as detailed in the [Prerequisites](#) and chose the fields available for configuration as detailed in the section [Selecting the Advanced Properties Displayed by Default](#).
- You must configure the properties in the addition/edition wizard of the object.

You can configure **DHCP servers, scopes and ranges** managing IPv4 with IPAM properties and DNS properties. For instance, you could create an IP address every time you add a DHCP static.

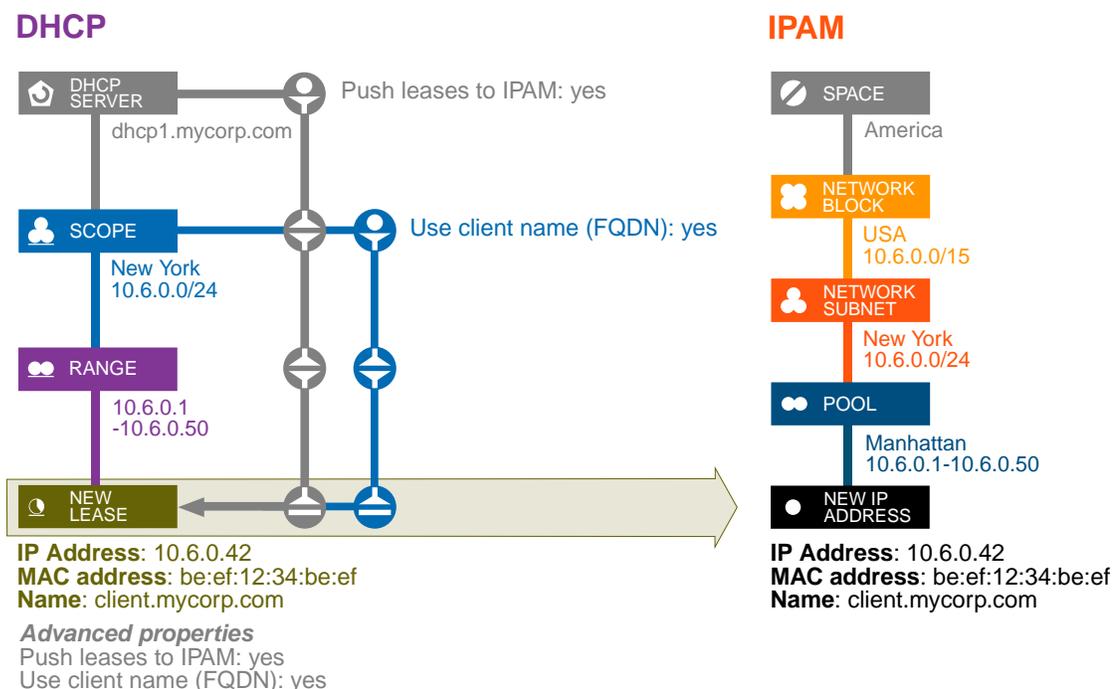


Figure 49.9. Replication from the DHCP to the IPAM at lease level

The table below details the entries of the customization wizard and the corresponding field in the addition/editing wizard.

Table 49.9. DHCP advanced properties fields

In the wizard Advanced properties customization	In the addition/editing wizard
Display the box "Push leases to IPAM"	Push leases to IPAM
Display the drop-down list "Lease name"	Lease name
Display the box "Use client name (FQDN)"	Use client name (FQDN)
Display the box "Update DNS"	Update DNS

All the properties detailed in the procedure below are based on a configuration that was not inherited. Most fields are already filled with the value in the container of the resource you create. For more details, refer to the chapter [Inheritance and Propagation](#).

To configure advanced properties on DHCP resources

1. In the sidebar, go to **DHCP > Servers, Scopes or Ranges** depending on your needs. The page opens.
2. On the right-end side of the menu, click on **[v4]**. The page refreshes and the button turns black.
3. Add or edit the DHCP object of your choice. The wizard opens. For more details, refer to the part [DHCP](#).
4. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to

be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on  to define its *Inheritance property* and/or *Propagation property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.10. DHCP advanced properties configuration fields

Field	Default value	Description
Push leases to IPAM	Unticked	Tick this box to assign an IP address in the IPAM for each allocated lease, using the client name and MAC address. If you set a space for the scope, the IP address is assigned or edited in that space. If you did not set any space for the scope, the IPAM update applies to any matching IP address in the smallest terminal network possible. Once you ticked this box, the fields <i>Lease name</i> , <i>Use client name (FQDN)</i> and <i>Update DNS</i> appear.
Lease name	Only one client can update the IPAM	Select the behavior of your choice in case of hostname conflict in the IPAM and DHCP. Such a conflict could affect the DNS, so you must decide what you want to send to the IPAM.
Only the first client updates the IPAM	/	If multiple clients that have the same name obtain a lease, only the client getting the first lease updates the IPAM and DNS. Until the first client's lease has not expired, their name cannot be replaced in the DNS. Basically, the first client getting a lease push their information in the IPAM (name, MAC address and IP address). The second client with the same name, has their IP and MAC addresses pushed to the IPAM without a name and therefore do not update the DNS.
Only one client can update the IPAM	/	This option is dedicated to mobile clients, on a network configuration composed of several scopes. It allows the client connecting from two different parts of the network with the same name and MAC address to be listed twice in the DHCP. The lease name, MAC address and IP address are pushed to the IPAM: the name and MAC address are identical but the IP address differs. Once saved in the IPAM, only the latest lease information updates the DNS.
Clients always update the IPAM	/	All the DHCP clients update the IPAM no matter their name, IP or MAC address. In other words, every client getting a lease has their information pushed to the IPAM (name, MAC address and IP address) even if their name has already been pushed to the IPAM. This is the most permissive mode. Keep in mind that only the latest lease information updates the DNS.
Use client name (FQDN)	Unticked	Tick this box to make sure the leases' "client name" value (the FQDN) is pushed in the IPAM and used as IP address name. If you tick the box <i>Update DNS</i> as well, the zone matching the FQDN is updated.
Update DNS	Unticked	Tick this box to update the DNS records database with the lease information.

5. Click on to complete the operation. The report opens and closes. The page is visible again.

Configuring DNS Advanced Properties

To configure advanced properties:

1. Your administrator must have configured the internal module setup as detailed in the [Prerequisites](#) and chose the fields available for configuration as detailed in the section [Selecting the Advanced Properties Displayed by Default](#).
2. You must configure the properties in the addition/edition wizard of the object.

You can configure **DNS servers, views and zones** with IPAM properties and DNS properties.

Enabling the advanced properties allows to automatically create an IP address when you add a DNS record.

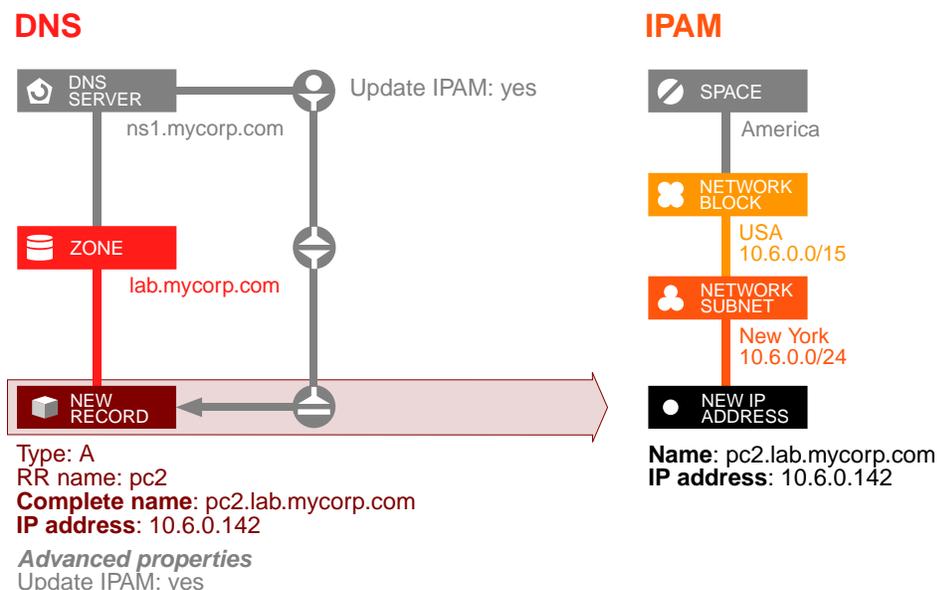


Figure 49.10. Replication from the DNS to the IPAM at record level

The table below details the entries of the customization wizard and the corresponding field in the addition/edition wizard.

Table 49.11. DNS advanced properties configuration fields

In the wizard Advanced properties customization	In the addition/edition wizard
Display the box "Update IPAM"	Update IPAM
Display the box "Create PTR" (not available for DNS views)	Create PTR

All the properties detailed in the procedure below are based on a configuration that was not inherited. Most fields are already filled with the value in the container of the resource you create. For more details, refer to the chapter [Inheritance and Propagation](#).

To configure advanced properties on DNS resources

1. Go to **DNS > Servers, Views or Zones**. The page opens.
2. Add or edit the DNS object of your choice. For more details, refer to the part [DNS](#). The corresponding wizard opens.
3. On the last page of the wizard, configure the advanced properties of your choice according to the table below:

If not all the fields are part of the wizard *Default* display, users with sufficient rights can select *All* in the drop-down list **Advanced properties** to display all the fields that do not require to be enabled in the customization wizard. All the fields appear but their order does not respect the table.

Next to each advanced property, you can click on  to define its *Inheritance property* and/or *Propagation property*. For more details, refer to the chapter [Inheritance and Propagation](#).

Table 49.12. DNS advanced properties configuration fields

Field	Default value	Description
Update IPAM	Unticked	Tick this box to update the IPAM when you create A, AAAA or PTR resource records. If you set a space for the zone, the IP address(es) is assigned or edited in that space. If you did not set any space for the zone, the IPAM update applies to any matching IP address in the smallest terminal network possible.
Create PTR	Unticked	Tick this box to create a PTR record for each A or AAAA resource record created in the DNS physical servers managed via a smart architecture. This option can be applied on DNS servers and zones. Note that if you create several A records with one name pointing to the several IP addresses, only the first PTR record is created.

- Click on to complete the operation. The report opens and closes. The page is visible again.

Setting Advanced Properties

Advanced properties can be set individually or several objects at a time regardless of any preexisting configuration. This operation is quite advanced and should not be delegated lightly to users that do not belong the group *admin*. Standard users can view but not modify the advanced properties parameters applied to an object. They appear in gray in the addition and edition wizards.

You cannot use this option on VLAN and Device Manager advanced properties.

Using the option *Set advanced properties* allows to:

- Set a property on an object.
- Overwrite any advanced property already configured.
- Overwrite the value of the *Inheritance property* and/or *Propagation property* of an advanced property.
- Propagate a configured advanced property to child objects if their *Inheritance property* is set to *Inherit*.
- Restrict an advanced property at the level of your choice.

Note that, in the wizard, the drop-down lists *Inheritance property* and *Propagation property* are only displayed if they are relevant to the object you selected.

To set advanced properties on several objects

This operation should only be performed by users of the group *admin*.

- Depending on your needs, in the sidebar:

- a. Go to **IPAM** > **Spaces, Networks, Pools** or **Addresses**. The page opens.
 - b. Go to **DHCP** > **Servers, Scopes** or **Ranges**. The page opens.
 - c. Go to **DNS** > **Servers, Views**, or **Zones**. The page opens.
2. Tick the object(s) for which you want to set advanced properties.
 3. In the menu, select **Tools** > **Expert** > **Set advanced properties**. The wizard **Set advanced properties** opens.
 4. In the drop-down list **Property**, select the advanced property of your choice. The page refreshes and the drop-down list **Value** appears.
 5. In the drop-down list **Inheritance property**, configure the advanced property inheritance behavior:
 - a. Select **Set** to set the advanced property or overwrite the value of the advanced property if it is already set on the object. This value is selected by default.
 - b. Select **Inherit** to inherit the advanced property from a parent object. Selecting this value hides the drop-down list *Value*.
 6. If you selected *Set*, in the drop-down list **Value** select the parameter value.
 7. In the drop-down list **Propagation property**, configure the advanced property propagation behavior:
 - a. Select **Propagate** to propagate the advanced property value on the child objects. The propagation is only possible if the child objects Inheritance property is *Inherit*. This value is selected by default.
 - b. Select **Restrict** to only configure the advanced property on the object(s) selected and prevent its propagation to the child objects.
 8. Click on **ADD** to move your advanced property configuration to the **Advanced properties list**.
 - To update an entry, select a configuration in the list, it is loaded again in the drop-down lists. Edit the needed data and click on **UPDATE**.
 - To delete an entry, select a configuration and click on **DELETE**.
 - To discard the latest modifications, click on **CANCEL**.
 9. Repeat the steps 5 to 9 for as many advanced properties as you need.
 10. Click on **OK** to complete the operation. The report opens and closes. Any property value and configuration previously set for the selected object(s) is overwritten.

Part IX. Application

The module Application allows to maintain an application inventory, tailor application traffic on your network and optimize user experience.

Every application you add is registered with a Fully Qualified Domain Name (FQDN) and configured with a set of pools and nodes that define its infrastructure. The application infrastructure allows to define a traffic policy.

Each traffic policy can be enforced on your network thanks to an existing DNS infrastructure. After enabling Global Server Load Balancing (GSLB) on compatible SOLIDserver appliances, you can associate application traffic policies with one or more GSLB servers and deploy the traffic policy on your network.

The way you configure your traffic policies can therefore maximize application availability and optimize the way users access your applications.

Once an application is deployed, your application traffic policy takes over the DNS resolution if and only if the query matches the FQDN of the application. If it does, the query is routed to the appropriate pool, the load balancing configuration of that pool determines which node is the most suited to answer the query and the node IP address is sent out to the querying client if its last health check was successful. Even if an existing record could have answered the query it does not, whether the server cached it or has authority over it, the standard DNS answer is ignored.

When traffic policy deployments are complete, the GSLB server only handles queries if they match a registered FQDN and if at least one application node is operational. In any other case, the DNS server answers the queries.

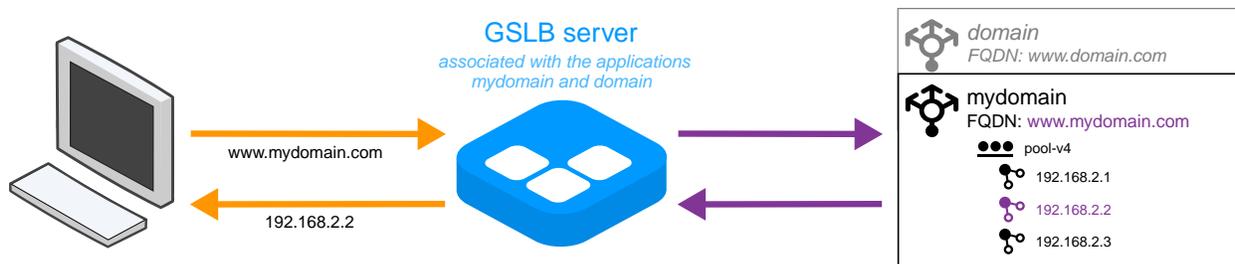


Figure 198. How a deployed application takes over the DNS resolution

The Application hierarchy includes 3 levels:

- **Applications:** The highest level of the hierarchy. It contains pools and nodes. An application is configured with a unique FQDN and once associated with one or more GSLB enabled server(s), it can deploy its configuration on your network. For more details, refer to the chapter [Managing Applications](#).
- **Pools:** The second level of the hierarchy. They contain nodes. A pool can manage IPv4 or IPv6 addressing, its load balancing configuration determines how the traffic is directed toward the nodes it manages. For more details, refer to the chapter [Managing Pools](#).
- **Nodes:** The last level of the hierarchy. They belong to a pool. They have a unique IPv4 or IPv6 address and can be configured with a health check that ensures they are responsive. For more details, refer to the chapter [Managing Nodes](#).

Before managing the Application objects you must configure the module as detailed in the chapter [Configuring Application](#).

Chapter 50. Configuring Application

To configure the module *Application* and use it to the fullest you must:

- Meet the [prerequisites](#).
- Take into account the [limitations](#).
- Configure and enable the [service GSLB Server](#).

Prerequisites

- A SOLIDserver appliance in version 7.1 or higher.
- An appliance with at least 8 GB of RAM:
 - Either one of our hardware appliance models, except SDS-260 and SDS-3300.
 - Or a virtual appliances must be configured with an *intel* network card (*em**, *ig**, *igb**, *ix**, *ixg**, *ixv** or *ixl**).
- The license DNS GSLB:
 - Without it you cannot associate applications with GSLB servers, and therefore you cannot deploy traffic policies on your network.
 - It must be activated on every SOLIDserver appliance meant to load balance application traffic.
- An appliance properly configured:
 - The communication between a GSLB enabled SOLIDserver and its clients has to be over UDP and/or TCP.
 - Both DNS and GSLB dedicated services must be running to ensure service continuity. Make sure that either the services *GSLB Server* and *DNS server (named)* or the services *GSLB Server* and *DNS server (unbound)* are both enabled and started.
 - Any FQDN registered for an application must be resolvable to be properly deployed. Either the GSLB server manages an A or AAAA record matching the FQDN or it can cache the initial FQDN answer.
 - Make sure traffic policies can actually take over the resolution. You must adapt the TTL of the resource records matching the FQDNs you register for any application. The new TTL must take into account your traffic policy details, the pool session duration you may configure and the configuration of your nodes health check.
- The appropriate rights and resources configured for end users.
 - All objects are visible to any user by default but end-users cannot manage them if they do not belong to a group granted access to all the module *Application* rights.
 - To be able to configure an application with a GSLB server, users must belong to a group that is granted the right *Display: DNS servers list* and has the server among its resources.

Limitations

- One DNS server cannot be used for DNSSEC and GSLB.

If you associate an application with a GSLB server used as DNSSEC resolver and/or that has signed zones, its resolution and/or answers are no longer DNSSEC compliant.

- DNS round-robin is not supported by the Application nodes.

If a node does not answer, the server associated with the application answers the query. However, any round-robin configuration in the DNS is ignored, only one record is used to answer the query, the rest of the records configured are ignored.

- Restoring a backup containing Application data overwrites the current configuration. Whether you created, edited or deleted objects since the backup was saved, all changes are lost and may not even be visible in the GUI. Before restoring a backup, make sure you saved it when the Application database was up-to-date.
- You can only configure and apply classes at application level. For more details regarding classes, refer to the chapter [Configuring Classes](#).
- The module *Dashboards* does not include a page dedicated to the module *Application*.

Configuring and Enabling the Service GSLB Server

Once you have met the [prerequisites](#) and have taken into account the [limitations](#), you must configure the listening interfaces of all GSLB enabled servers and enable the service. Without this configuration, the traffic policies are listed in the module but never deployed on your network.

The DNS must be running as well. Make sure the services *GSLB server* and *DNS server* (*named* / *nsd* / *unbound*) are both enabled and started.

Note that if your license includes both DNS Guardian and DNS GSLB, you must configure the line *DNS Guardian / GSLB server* as both features rely on the same service.

To configure the listening interfaces of GSLB server and enable the service

Only users of the group *admin* can perform this operation.

1. Configure the listening interfaces of GSLB server

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
- c. In the column **Name**, click on **GSLB server** or **DNS Guardian / GSLB server**. The wizard **GSLB server configuration** or **DNS Guardian & GSLB server configuration** opens.
- d. In the list **Available interfaces**, select the interface of your choice and click on . The interface is moved to the list **Selected interfaces**.

Each interface is listed *<interface-name>* (*<MAC-address>*), whether it is active or not. Only *Intel* network interfaces are listed as no other interface card can be configured for the service.

- e. Repeat this action for as many interfaces as you need.

To remove an interface from the list **Selected interfaces**, select it and click on . The interface is moved back to the list **Available interfaces**.

- f. Click on  to complete the operation. The report opens and closes.

2. **Enable the service GSLB server**

- a. In the column **Name**, look for *GSLB server* or *DNS Guardian / GSLB server*.
- b. In the column **Enabled**, click on the link **Disabled** to enable the service. The wizard opens.
- c. Click on to complete the operation. The wizard closes. The page is visible again.

3. **Apply your configuration**

- a. Right now your configuration is pending. In the menu, select . **Tools > Apply configuration** to save your changes. The wizard **Commit the system configuration changes** opens.
- b. Click on to complete the operation. The page refreshes.

If the service is marked  *Warning*, it might mean that the system has not enough memory installed, that a configuration file is corrupt or that the license is not valid.

Once the service is enabled, on the page *All servers* all compatible DNS servers are marked  *Enabled* in the column **Guardian/GSLB**.

To stop or disable the service *GSLB server*, refer to the section [Handling Services](#) in the chapter *Configuring the Services*.

Chapter 51. Managing Applications

From the page *All applications* you can manage applications and set up entire traffic policies.

Each application must be configured with a Fully Qualified Domain Name (FQDN) that can be resolved.

To deploy an application or traffic policy on your network you must associate it with one or more GSLB enabled servers able to resolve the application FQDN:

- On recursive servers, the initial query of the FQDN is cached in order to answer all the following queries.
- On authoritative server, the FQDN must be declared as A or AAAA record in one of its zones.

Once an application is deployed on at least one GSLB server, when a client queries a domain:

1. The application handles the query if it is configured with the relevant FQDN.
2. It directs the query to the relevant pool and node.

You can either create an application and then create the pools and nodes it contains, or you can create the application and its traffic policy from the page *All applications*. For more details, refer to the section [Adding and Deploying Applications and Traffic Policies](#).

Browsing Applications

The application is the highest level of the Application hierarchy.

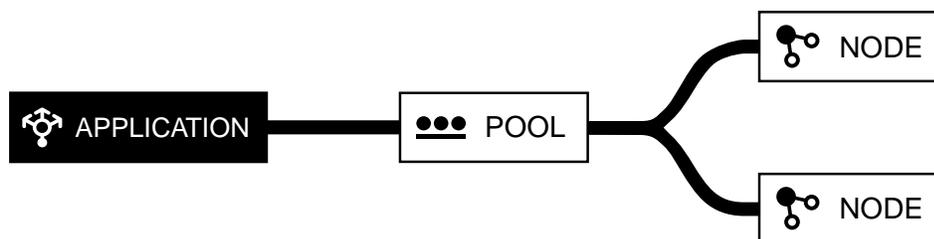


Figure 51.1. The application in the Application hierarchy

Browsing the Application Database

To display the list of applications

1. In the sidebar, go to **Application** > **Applications**. The page **All applications** opens.
2. To display or hide deployed applications, click on . The page refreshes.

In the column **Name**, the icon precedes every application.

If there are deployed applications, they are preceded by the icon and listed under the parent application name as many times as there are GSLB servers associated with the application.

To display an application properties page

1. In the sidebar, go to  **Application** > **Applications**. The page **All applications** opens.
2. At the end of the line of the application of your choice, deployed or not, click on . The application properties page opens.

Customizing the Display on the Page All Applications

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

All the columns are displayed by default.

Table 51.1. The columns on the page All applications

Column	Description
Name	The name of the application.
FQDN	The FQDN of the application.
GSLB server	The name of the physical server associated with a deployed application, it answers the queries made to the application FQDN.
Class	The class applied to the application, the column is empty for the deployed application and traffic policies. For more details, refer to chapter Configuring Classes .
Multi-status	Status details on the deployed nodes of the application. The column is always empty for applications not associated with any server.
Status	The status of the application or the service <i>GSLB server</i> . For more details, refer to the section Understanding the Statuses on the Page All Applications .

Understanding the Statuses on the Page All Applications

The column **Status** returns information regarding the applications, their content and the service *GSLB server*.

Table 51.2. The statuses on the page All applications

Status	Description
 OK	The application and its content are configured and running.
 Warning	The <i>Operational status</i> of at least one of the nodes of the application is <i>Inactive</i> .
 Down	The <i>Operational status</i> of all the nodes of the application is <i>Inactive</i> .
 Delayed delete	The content of the application is being deleted from the physical servers it is deployed on.
 GSLB Timeout	The service <i>GSLB server</i> is unreachable.
 GSLB Invalid credentials	The SSL credentials of the <i>GSLB server</i> associated with the application are invalid or this server is already managed by another appliance and you need to specify your credentials again. For more details, refer to the section Editing DNS Servers .
 GSLB Stopped	The service <i>GSLB server</i> is not running.

Adding and Deploying Applications

From the page *All applications* you can add applications, they are the entry point to managing pools and nodes. To add an application and its content in a single wizard, refer to the section [Adding and Deploying Applications and Traffic Policies](#).

Before adding an application, keep in mind that:

- An application must have a unique name and FQDN.

Several applications can share the same name if they manage different FQDNs. In the same way, several applications can manage the same FQDN if they have different names.

- The FQDN of an application must be resolvable.
- To deploy an application and its content on your network, you must associate it with at least one GSLB server.
- A GSLB server can only be associated once with one FQDN.

Once associated with an application, a GSLB server cannot be associated with another application managing the same FQDN.

- Any FQDN registered for an application must be resolvable to be properly deployed. Either the GSLB server manages an A or AAAA record matching the FQDN or it can cache the initial FQDN answer.

You can add as many applications as you need.

To add an application

1. In the sidebar, go to  **Application > Applications**. The page **All applications** opens.
2. In the menu, select **+ Add > Application**. The wizard **Add an application** opens.
3. If you or your administrator created application classes, in the list **Application class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the field **Name**, type in the name of the application.
5. In the field **FQDN**, type in the Fully Qualified Domain Name of the application.
6. In the list **Available GSLB servers**, you can select a server and click on . The server is moved to the list **Selected GSLB servers**.

Only GSLB enabled physical servers are listed, whether they are managed by a smart architecture or not.

To remove a server from the **Selected GSLB servers**, select it and click on . The server is moved back to the list **Available GSLB servers**.

7. Repeat the operation for as many as servers as needed.
8. Click on **OK** to complete the operation. The report opens and closes. The application is listed.

If you associated the application with one or more GSLB servers, click on . Several lines appear under the application itself, there is a line for each of the server(s) the application is deployed on.

Once added, you can edit the GSLB server association of an application. For more details, refer to the section [Editing Applications](#).

Adding and Deploying Applications and Traffic Policies

From the page *All applications* you can add applications and their traffic policies, that is to say a pool and its nodes, in one wizard.

Before adding an application and its traffic policy keep in mind that:

- An application must have a unique name and FQDN.
Several applications can share the same name if they manage different FQDNs. In the same way, several applications can manage the same FQDN if they have different names.
- To deploy an application and its traffic policy on your network, you must associate it with at least one GSLB server.
- Any FQDN registered for an application must be resolvable to be properly deployed. Either the GSLB server manages an A or AAAA record matching the FQDN or it can cache the initial FQDN answer.
- A GSLB server can only be associated once with one FQDN.

Once associated with an application, a GSLB server cannot be associated with another application managing the same FQDN.

- After adding an application and traffic policies, the application, pool and node(s) it contains are managed individually. To edit any, you must refer to the edition section of each object.

To add an application

1. In the sidebar, go to  **Application > Applications**. The page **All applications** opens.
2. In the menu, select **+ Add > Application and traffic policy**. The wizard opens.
3. If you or your administrator created application classes, in the list **Application class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the field **Name**, type in the name of the application.
5. In the field **FQDN**, type in the Fully Qualified Domain Name of the application.
6. In the list **Available GSLB servers**, you can select a server and click on **+**. The server is moved to the list **Selected GSLB servers**.

Only GSLB enabled physical servers are listed, whether they are managed by a smart architecture or not.

To remove a server from the **Selected GSLB servers**, select it and click on **-**. The server is moved back to the list **Available GSLB servers**.

7. Repeat the operation for as many as servers as needed.
8. Click on **NEXT**. The page **Add a pool** opens.
9. Configure the pool:

Field	Description
Name	The name of the pool. This field is required.

Field	Description
Protocol	The version of the IP protocol of pool and its nodes: <i>IPv4</i> or <i>IPv6</i> . This field is required.
Load balancing mode	The mode of the pool: <i>Round-robin</i> , <i>Latency</i> or <i>Weighted</i> . This field is required.
<i>Round-robin</i>	The traffic is evenly directed to all active nodes. This mode is selected by default.
<i>Latency</i>	The traffic is directed to the active nodes with the best latency. If you select this mode, the field <i>Max. preferred nodes</i> appears. <i>Max. preferred nodes</i> : The number of active nodes with the best latency that must answer queries.
<i>Weighted</i>	The traffic is directed to the active nodes depending on their weight.
Enable session affinity	Tick the box if you want to set a period of time during which all traffic is directed to the nodes. If you tick the box, the field <i>Session duration</i> appears.
<i>Session duration</i>	Specify the period of your choice, in seconds. By default, it is set to <i>300</i> . This field is required.

10. Click on **[NEXT]**. The page **Add a node** opens.

11. Create at least one node.

a. Configure the node:

Field	Description
Name	The name of the node. Each node name must be unique.
IP address	Type in the IPv4 or IPv6 address of the node, depending on the selected <i>Pool protocol</i> . Each node IP address must be unique.
Mode	The mode of the node, either <i>Active</i> or <i>Backup</i> . If the pool is configured with the load balancing mode <i>Latency</i> , the field is in read-only and the node is <i>Active</i> by default.
Health Check type	Select <i>None</i> , <i>HTTP(S)</i> , <i>Ok</i> , <i>Ping</i> or <i>TCP</i> . <i>None</i> is selected by default.
<i>None</i>	There is no health check configured for the node, whether the node is answering queries or not, its <i>Status</i> is always <i>OK</i> .
<i>HTTP(S)</i>	The health check of the node is performed via HTTPS.
<i>Ok</i>	The health check of the node always marks the node <i>Operational status</i> as <i>OK</i> .
<i>Ping</i>	The health check of the node is performed using ping commands.
<i>TCP</i>	The health check of the node is performed via TCP.

b. Click on **[ADD]**. The node is moved to the **Node list** and listed as follows: *<node-name> - <IP-address> - <health-check-type> (<mode>)*.

- To update a node, select it in the list. Its information is loaded in the fields, change data according to your needs and click on **[UPDATE]**.
- To delete a node, select it in the list and click on **[DELETE]**.
- To discard the latest changes, click on **[CANCEL]**.

c. Repeat the operation for as many nodes as needed.

12. Click on **[OK]** to complete the operation. The report opens and closes. The application is listed.

If you associated the application with one or more GSLB servers, click on . Several lines appear under the application itself, there is a line for each of the server(s) the application is deployed on.

If you associated the application with a GSLB server and configured it with a resolvable FQDN, the application and the traffic policies are deployed right away.

Once added, you cannot edit a traffic policy. However you can individually edit the application, pool and node(s) that compose it. For more details, refer to the sections [Editing Applications](#), [Editing Pools](#) and [Editing Nodes](#).

Editing Applications

At any time you can edit applications, whether they were created via the addition menu *Application* or *Application and traffic policy*.

Before editing an application, note that:

- You cannot edit the name or FQDN of an application.
- You cannot edit a  deployed application. You must edit the application itself.
- The application edition consists of:
 - Associating it with extra GSLB servers or different ones.
 - Dissociating it from one or more GSLB servers.

To edit an application

1. In the sidebar, go to  **Application > Applications**. The page **All applications** opens.
2. Filter the list if need be.
3. At the end of the line of the application of your choice, deployed or not, click on . The application properties page opens.
4. In the panel **Main properties**, click on **EDIT**. The wizard **Edit an application** opens.
5. If you or your administrator created application classes, in the list **Application class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. The **Name** and **FQDN** are displayed in a read-only gray field. You cannot edit them.
7. Edit the list **Selected GSLB servers** according to your needs. Only GSLB enabled physical servers are listed, whether they are managed by a smart architecture or not.

Select a server in the list *Available GSLB servers* and click on  to add it the list *Selected GSLB servers*. Select a server in the list *Selected GSLB servers* and click on  to remove it move it back to the list *Available GSLB servers*.

8. Repeat the operation for as many as servers as needed.
9. Click on **OK** to complete the operation. The report opens and closes. The application properties are updated.

Deleting Applications

At any time you can delete an application. Note that:

- Deleting an application also deletes the pools and nodes it contains.
- You cannot delete a  deployed application.

You must either delete the application itself or dissociate the relevant GSLB server from the application, the dedicated deployment line is no longer listed. For more details, refer to the section [Editing Applications](#).

To delete an application

1. In the sidebar, go to  **Application > Applications**. The page **All applications** opens.
2. Filter the list if need be.
3. Tick the application(s) you want to delete.
4. In the menu, click on . The wizard **Delete** opens.
5. Click on  to complete the operation. The report opens and closes. The application is no longer listed, the deployed application lines are deleted as well.

Chapter 52. Managing Pools

From the page *All pools* you can manage IPv4 and IPv6 pools. Pools belong to applications and manage nodes.

If the application they belong to is associated with a GSLB enabled server, the pools are deployed on your network.

When an application is deployed on at least one GSLB server, when a client queries a domain:

1. The application handles the query if it is configured with the relevant FQDN.
2. It directs the query to the appropriate pool.
3. The pool load balancing configuration determines towards which node the query is directed.

Browsing Pools

The pool is the second level of the Application hierarchy.

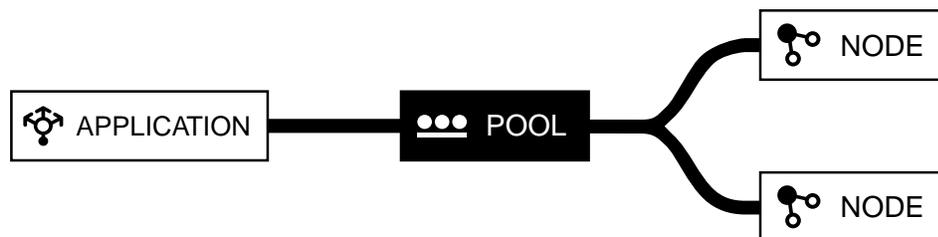


Figure 52.1. The pool in the Application hierarchy

Browsing the Pool Database

To display the list of pools

1. In the sidebar, go to **Application** > **Pools**. The page **All pools** opens.
2. If need be, filter the column **Protocol** to display only IPv4 or IPv6 pools.
3. To display or hide the pools of deployed applications, click on . The page refreshes.

The same node name can be listed multiple times. In the column **Application name**, each application is listed once, preceded by and, if it is deployed, its name is preceded by and is listed as many times as their are GSLB server associated with it.

4. To display the pools of a specific application, in the column **Application Name**, click on the name of the application of your choice. The page refreshes.

To display a pool properties page

1. In the sidebar, go to **Application** > **Pools**. The page **All pools** opens.
2. At the end of the line of the pool of your choice, whether it belongs to a deployed application or not, click on . The pool properties page opens.

Customizing the Display on the Page All Pools

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

All the columns are displayed by default.

Table 52.1. Available columns on the page All applications

Column	Description
Name	The name of the pool.
Protocol	The protocol version of the pool, either <i>IPv4</i> or <i>IPv6</i> .
Mode	The load balancing mode of the pool, either <i>Round-robin</i> , <i>Latency</i> or <i>Weighed</i> .
Nodes (max.)	The maximum number of active nodes with the lowest latency of the pool that answer the queries made to the application FQDN. It only contains data for pools configured with the load balancing mode <i>Latency</i> .
Application name	The name of the application the pool belongs to.
Application FQDN	The FQDN of the application the pool belongs to.
Application GSLB	The name of the physical server associated with the deployed application the pool belongs to.
Affinity state	The pool session affinity state, either <i>Enabled</i> or <i>Disabled</i> .
Session duration (sec.)	The affinity state session duration, in seconds.
Multi-status	Status details on the deployed nodes of the pool. The column is always empty for applications not associated with any server.
Status	The status of the pool, either <i>Delayed create</i> , <i>Delayed delete</i> , <i>Warning</i> , <i>Down</i> or <i>OK</i> .

Understanding the Statuses on the Page All Pools

The column **Status** returns information regarding the pools you manage and their content.

Table 52.2. The statuses on the page All pools

Status	Description
 Delayed create	The content of the pool is being deployed on the physical servers associated with the application it belongs to.
 Delayed delete	The content of the pool is being deleted from the physical servers associated with the application it belongs to.
 Warning	The <i>Operational status</i> of at least one of the nodes of the pool is <i>Inactive</i> .
 Down	The <i>Operational status</i> of all the nodes of the pool is <i>Inactive</i> .
 OK	The nodes of the pool are configured and running.

Adding Pools

From the page *All pools* you can add IPv4 and IPv6 pools in your applications, they allow to manage nodes.

The page may list pools created from the page *All applications* via the menu *Application and traffic policy*. For more details, refer to the section [Adding and Deploying Applications and Traffic Policies](#).

Before adding a pool, keep in mind that:

- A pool must belong to an application.
- Each pool must have a unique name in one application.
- You cannot add more than two pools in one application.
- You cannot manage more than one IPv4 pool and one IPv6 pool in one application.

To add a pool

1. In the sidebar, go to  **Application > Pools**. The page **All pools** opens.
2. In the menu, select **+ Add**. The wizard **Add a pool** opens.
3. If you or your administrator created application classes, in the list **Application class** select one class, *All* or *No class*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the list **Application name**, select the application of your choice. If some application classes are enabled, only the applications matching the selected *Application class* are displayed.

Click on **NEXT**. The last page of the wizard opens.

5. In the field **Name**, type in the name of the pool.
6. In the drop-down list **Protocol**, select *IPv4* or *IPv6*. The page refreshes. By default, *IPv4* is selected.
7. In the drop-down list **Load balancing mode**, select one of the following modes. The page refreshes.

Mode	Description
Round-robin	The traffic is evenly directed to all active nodes. Within the pool, all active nodes weight <i>1</i> and backup nodes weight <i>0</i> . This mode is selected by default.
Latency	The traffic is directed to the active nodes with the best latency. If you select this mode, the field <i>Max. preferred nodes</i> appears. <i>Max. preferred nodes</i> : Specify the maximum number of nodes with the lowest latency that must answer the queries made to the application FQDN. Only the active nodes of the pool answer the queries. By default, it is set to <i>1</i> .
Weighed	The redirection of traffic depends on the weight you set for the active nodes, the nodes set with the greater weight answer first. All backup nodes weight <i>0</i> .

8. You can tick the box **Enable session affinity** to set a period of time during which the pool sends out the same answer to a given client, no matter how many times they query the same information. If you tick the box, the field *Session duration* appears.
9. In the field **Session duration**, specify the duration of the session affinity, in seconds. By default, it is set to *300*.
10. Click on **OK** to complete the operation. The report opens and closes. The pool is listed.

If the application it belongs to is associated with one or more GSLB servers, click on . Several lines appear under the pool itself, there is a line for each of the server(s) the application is deployed on.

Editing Pools

At any time you can edit pools, whether they were created from the page *All pools* or the page *All applications* via the menu *Application and traffic policy*.

Before editing a pool, note that:

- You cannot edit a pool protocol.
- You cannot edit a deployed pool. You must edit the corresponding pool.

In the column *Application name*, a deployed pool belongs to a  deployed application.

To edit a pool

1. In the sidebar, go to  **Application > Pools**. The page **All pools** opens.
2. Filter the list if need be.
3. At the end of the line of the pool of your choice, whether it belongs to a deployed application or not, click on . The pool properties page opens.
4. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a pool** opens.
5. If you or your administrator created application classes, in the list **Application class** select one class, *All* or *No class*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Edit the pool **Name**, **Load balancing mode** and/or **Enable session affinity** configuration according to your needs.
7. Click on **[OK]** to complete the operation. The report opens and closes. The pool properties are updated.

Deleting Pools

At any time you can delete a pool. Note that:

- Deleting a pool also deletes the nodes it contains.
- You cannot delete a deployed pool. In the column *Application name*, a deployed pool belongs to a  deployed application.

You must either delete the pool itself or dissociate the relevant GSLB server from the parent application. Once the GSLB server is dissociated from the application, the dedicated pool deployment line is no longer listed. For more details, refer to the section [Editing Applications](#).

To delete a pool

1. In the sidebar, go to  **Application > Pools**. The page **All pools** opens.
2. Filter the list if need be.
3. Tick the pool(s) you want to delete.
4. In the menu, click on . The wizard **Delete** opens.

5. Click on to complete the operation. The report opens and closes. The pool is no longer listed, the deployed pool lines are deleted as well.

Chapter 53. Managing Nodes

From the page *All nodes* you can manage IPv4 and IPv6 nodes. As nodes belong to pools they are configured with either an IPv4 or IPv6 address.

If the application they belong to is associated with a GSLB enabled server, the nodes are deployed on your network.

A node is the endpoint of a query. Each node must be configured with a unique IP address and can be configured with a health check that verifies its status allows it to answer queries.

When a traffic policy is set and deployed on at least one GSLB server, when a client queries a domain:

1. The application handles the query if it is configured with the relevant FQDN.
2. It directs the query to the appropriate pool.
3. All the node health checks are executed within the pool. If the health checks succeed and the node status is cleared, the IP address of one or all the nodes can be sent out to the client.
4. Depending on the pool load balancing mode and configuration, it determines which node is best suited to answer the query.

Browsing Nodes

The node is the lowest level of the Application hierarchy.

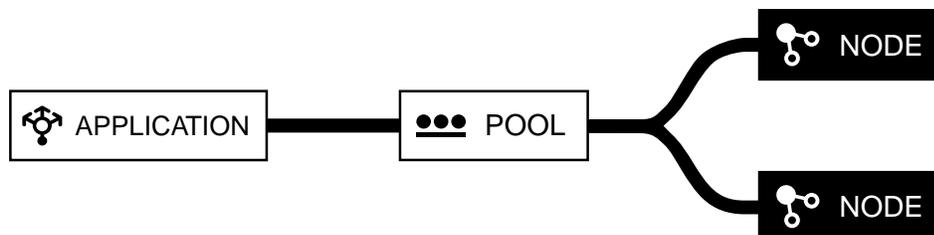


Figure 53.1. The node in the Application hierarchy

Browsing the Node Database

To display the list of nodes

1. In the sidebar, go to **Application > Nodes**. The page **All nodes** opens.
2. You can filter the columns **IPv4 address** and **IPv6 address** to only display IPv4 or IPv6 nodes.
3. To display or hide the nodes of deployed applications, click on **[icon]**. The page refreshes.

The same node name can be listed multiple times. In the column **Application name**, each application is listed once, preceded by **[icon]** and, if it is deployed, its name is preceded by **[icon]** and is listed as many times as their are GSLB server associated with it.

4. To display the nodes of a specific application, in the column **Application Name**, click on the name of the application of your choice. The page refreshes.

- To display the nodes of a specific pool, in the column **Pool Name**, click on the name of the pool of your choice. The page refreshes.

Note that for IPv6 addresses, you can compress or uncompress the display via the button **0::0** or display IPv6 labels above parts of the addresses listed via the button **lab:0**. For more details, refer to the chapter [Managing IPv6 Labels](#).

To display a node properties page

- In the sidebar, go to **Application > Pools**. The page **All pools** opens.
- At the end of the line of the node of your choice, whether it belongs to a deployed application or not, click on **⌵**. The node properties page opens.

Customizing the Display on the Page All Nodes

Users of the group *admin* can create customized column layouts. The button **⌵ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

The following columns are displayed by default.

Table 53.1. Default columns on the page All nodes

Column	Description
Name	The name of the node.
IPv4 address	The IP address of the node if it belongs to an IPv4 pool. The column displays <i>N/A</i> if the pool belongs to an IPv6 pool.
IPv6 address	The IP address of the node if it belongs to an IPv6 pool. The column displays <i>N/A</i> if the pool belongs to an IPv4 pool.
Pool name	The name of the pool the node belongs to.
Application name	The name of the application the node belongs to.
Application FQDN	The FQDN of the application the node belongs to.
Application GSLB	The name of the physical server associated with the deployed application the node belongs to.
Mode	The load balancing mode of the pool the node belongs to, either <i>Round-robin</i> , <i>Latency</i> or <i>Weighed</i> .
Weight	The weight of the node, backup nodes are set to <i>0</i> , any other weight means that the node is active.
Health check (HC)	The selected health check of the node, either <i>None</i> , <i>Custom</i> , <i>HTTP(S)</i> , <i>Ok</i> , <i>Ping</i> or <i>TCP</i> .
Operational status	The operating status of a deployed node. For more details, refer to the table node operating statuses .
Status	The management status of each node. For more details, refer to the table node statuses .

Extra columns, dedicated to health checks are available. They display the *Expert mode* configuration details or *N/A*.

Table 53.2. Available columns on the page All nodes

Column	Description
HC Parameters	Any extra parameter you configured for the health check.
HC Failback threshold	The number of times, between <i>1</i> and <i>10</i> , the health check should return the same result before setting the node <i>Operational status</i> to <i>Active</i> .

Column	Description
HC Failover threshold	The number of times, between 1 and 10, the health check should return the same result before setting the node <i>Operational status</i> to <i>Inactive</i> .
HC Frequency	The number of seconds between health checks, either 10, 30, 60 or 500. The total number of health checks performed depend on the <i>HC Failback threshold</i> and <i>HC Failover threshold</i> .
HC Timeout	The health check timeout, between 1 and 10 seconds. Beyond this period, the health check times out if the node is not responding.

Understanding the Statuses on the Page All Nodes

The column **Operational status** provides status information on deployed nodes. The health check results determine this status.

Table 53.3. The node operating statuses

Status	Description
 Inactive	The health check configured for the node determined that the node is down.
 Active	The health check configured for the node determined that the node is up.
 Unknown	The operating status of the node at first, it changes after the initial health check.
 N/A	The operating status is not applicable. The column only displays values for deployed nodes.

The column **Status** provides management status information on all nodes.

Table 53.4. The node statuses

Status	Description
 Delayed create	The content of the node is being deployed on the physical servers associated with the application it belongs to.
 Delayed delete	The content of the node is being deleted from the physical servers associated with the application it belongs to.
 Unmanaged	The node is unmanaged.
 OK	The node is managed, configured and running.

Adding Nodes

From the page *All nodes* you can add nodes to IPv4 and IPv6 pools.

The page may list nodes created from the page *All applications* via the menu *Application and traffic policy*. For more details, refer to the section [Adding and Deploying Applications and Traffic Policies](#).

Before adding a node, keep in mind that:

- A node must belong to a pool.
- Each node must have a unique name in one pool.
- Each node must have a unique IP address in one pool.
- Each node can be configured with a health check.
- The health checks can only be performed if the GSLB server(s) associated with the application the node belongs to and the node itself can communicate.

- Within a pool configured with the load balancing mode *Latency*:
 - All nodes must be configured with the same health check type to guaranty a fair comparison of the latency of each node and ensure an appropriate traffic redirection.
 - You must not configure nodes with the health check *None* because any node configured without health check is directed traffic, in addition to the configured Maximum number of preferred nodes.

Note that you can display labels on IPV6 nodes, for more details refer to the chapter [Managing IPv6 Labels](#).

To add a node

1. In the sidebar, go to **Application > Nodes**. The page **All nodes** opens.
2. In the menu, select **+ Add**. The wizard **Add a node** opens.
3. If you or your administrator created application classes, in the list **Application class** select one class, *All* or *No class*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the list **Application name**, select the application of your choice, each application is followed by its FQDN. If application classes are enabled, only the application(s) matching the selected *Application class* are displayed.

Click on **NEXT**. The next page of the wizard opens.

5. In the list **Pool name**, select the pool of your choice, each pool is listed as follows: *<pool-name> [protocol-version]*.

Click on **NEXT**. The last page of the wizard opens.

6. In the field **Name**, type in the name of the node.
7. In the field **IP address**, type in the IPv4 or IPv6 address of the node, depending on the selected *Pool protocol*.
8. In the drop-down list **Mode**, select *Active* or *Backup*.

This field is not displayed if you selected a pool with the load balancing mode *Latency*, as all nodes are active by default.

9. In the field **Weight**, specify the value of your choice. It must be an integer between 0 and 255. Within the pool, the active nodes with the greater weight are queried first.

This field is only displayed for *Active* nodes belonging to a pool set with the load balancing mode *Weighted*. All *Backup* nodes have a weight of 0.

10. In the drop-down list **Health Check type**, select one of the following types. The page refreshes. For all types, except *None*, one or more fields appear.

Health Check type	Description
None	No health check is executed on the node, no further configuration is needed, you can go straight to the last step. <i>None</i> is the default value. You should not select this type if the node belongs to a pool with the load balancing mode <i>Latency</i> .
Custom	The health check is performed by the script of your choice.

Health Check type	Description
HTTP(S)	The health check is performed via HTTP or HTTPS.
Ok	The health check always returns its <i>Operational status</i> as <i>OK</i> .
Ping	The health check ensures that the node IP address answers a ping command.
TCP	The health check ensures that the TCP port is listening.

11. Configure the health check following the [table](#) below.
12. Click on to complete the operation. The report opens and closes. The node is listed.

If the application it belongs to is associated with one or more GSLB servers, click on . Several lines appear under the node itself, there is a line for each of the server(s) the application is deployed on.

Table 53.5. The health check configuration and expert parameters

Field	Description
For all Health Check types	
Expert mode	You can tick the box to configure the health check details. The fields <i>Health check frequency</i> , <i>Health check failover threshold</i> , <i>Health check fallback threshold</i> and/or <i>Health check timeout in seconds</i> appear.
<i>Health check frequency</i>	You can specify the frequency of the node health check, either <i>every 10 seconds</i> , <i>every 30 seconds</i> , <i>every minute</i> or <i>every 5 minutes</i> . By default, it is set to <i>every 30 seconds</i> . If you selected the Health Check type <i>Ok</i> , this is the only field available.
<i>Health check failover threshold</i>	Specify the number of times, between <i>1</i> and <i>10</i> , before the node <i>Operational status</i> is set to <i>Inactive</i> due to the health check result. By default, it is set to <i>3</i> . This field is required.
<i>Health check fallback threshold</i>	Specify the number of times, between <i>1</i> and <i>10</i> , before the node <i>Operational status</i> is set to <i>Active</i> due to the health check result. By default, it is set to <i>3</i> . This field is required.
<i>Health check timeout in seconds</i>	Specify the number of seconds, between <i>1</i> and <i>10</i> , before the health check times out if the node is not responding. This field is required.
Extra fields for the Health Check type Custom	
Script name	Specify the name of the script performing the health check of the node. Only the name of the file is accepted, do not specify its extension <i>fct_</i> prefix. The file must be located in the directory <i>/data1/gslb/script</i> of the appliance. This field is required.
Expert mode	You can tick this box to configure the <i>Custom</i> health check details. The field <i>Script parameters</i> appears.
<i>Script parameters</i>	You can specify the parameters of your script.
Extra fields for the Health Check type HTTP(S)	
Expert mode	You can tick this box to configure the <i>HTTP(S)</i> health check conditions. The fields <i>HTTP host</i> , <i>HTTP port</i> , <i>URL to retrieve</i> , <i>Use HTTPS</i> , <i>Check HTTPS certificate</i> , <i>Expected HTTP status</i> , <i>Response string to match</i> and <i>HTTP authentication credentials</i> appear.
<i>HTTP host</i>	You can specify the name of the virtual host providing the Server Name Indication (SNI) during the HTTP checks. The virtual host must be associated with application. If you do not specify a host, the IP address is used during the check.
<i>HTTP port</i>	Specify the port to use during the health check. It must be an integer greater than <i>0</i> . If you do not specify any port, the port <i>80</i> is used for a health check via HTTP and the port <i>443</i> is used for a health check via HTTPS.
<i>URL to retrieve</i>	You can specify the end of the URL that the node should respond to. The full URL to retrieve and check is composed of the specified <i>HTTP host</i> , <i>HTTP port</i> and <i>URL to retrieve</i> as follows: <i><http-or-https>://<http-host>:<http-port>/<URL-to-retrieve></i> . If you leave the field empty, <i>/</i> is used to complete the URL.

Field	Description
<i>Use HTTPS</i>	Select <i>Yes</i> or <i>No</i> . Either the protocol <i>HTTP</i> or <i>HTTPS</i> precedes the full URL of the health check. If this field is set to <i>Yes</i> , the field <i>Check HTTPS certificate</i> is displayed. By default, it is set to <i>Yes</i> . This field is required.
<i>Check HTTPS certificate</i>	Select <i>Yes</i> or <i>No</i> . If set to <i>Yes</i> , it verifies the validity of the SSL certificate of the specified <i>HTTP host</i> , if it is not valid the health check fails. By default, it is set to <i>No</i> . This field is only displayed if <i>Use HTTPS</i> is set to <i>Yes</i> . This field is required.
<i>Expected HTTP status</i>	You can specify one or more HTTP status codes to match in answer, if any other code is returned the health check fails. The field accepts an integer composed of up to 3 digits between 1 and 999, or a range of codes using x, e.g. 3xx matches any code between 300 and 399. If you leave the field empty, the HTTP code returned is ignored during the health check.
<i>Response string to match</i>	You can specify a string to match in the body of the answer. If the string specified is not matched at least partially in the body of the answer, the health check fails. If you leave the field empty, the body of the answer is ignored during the health check.
<i>HTTP authentication credentials</i>	You can specify the authentication credentials used to perform the health check following the format <code><login>:<password></code> .
Extra field for the Health Check type TCP	
<i>Expert mode</i>	You can tick this box to configure the <i>TCP</i> health check details. The field <i>TCP port</i> appears.
<i>TCP port</i>	Specify the port to use during the health check. It must be an integer greater than 0. If you do not specify any port, the port 80 is used.

Editing Nodes

At any time you can edit nodes, whether they were created from the page *All nodes* or the page *All applications* via the menu *Application and traffic policy*.

Before editing a node, note that:

- You cannot edit the IP address of a node.
- You cannot edit a deployed node. You must edit the corresponding node.

In the column *Application name*, a deployed node belongs to a  deployed application.

To edit a node

1. In the sidebar, go to  **Application > Nodes**. The page **All nodes** opens.
2. Filter the list if need be.
3. At the end of the line of the node of your choice, whether it belongs to a deployed application or not, click on . The node properties page opens.
4. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a node** opens.
5. If you or your administrator created application classes, in the list **Application class** select one class, *All* or *No class*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Edit the node **Name**, **Mode**, **Health Check type** and/or **Expert mode** configuration according to your needs. For more details regarding the types and expert configuration, refer to the [table](#) above.

7. Click on **OK** to complete the operation. The report opens and closes. The pool properties are updated.

Managing or Unmanaging Nodes

At any time you can stop managing nodes, or manage them again.

By default, all nodes are managed when you create them. Unmanaging a node allows to make sure that queries are never redirected toward the node you chose.

Keep in mind that you cannot unmanage a deployed node. In the column *Application name*, a deployed node belongs to a  deployed application.

To manage/unmanage a node

1. In the sidebar, go to  **Application > Nodes**. The page **All nodes** opens.
2. Filter the list if need be.
3. Tick the node(s) of your choice.
4. In the menu, select  **Edit > Manage > Yes** or **No**. The wizard opens.
5. Click on **OK** to complete the operation. The report opens and closes. In the column **Status**, the selected nodes are marked  **OK** or  **Unmanaged**.

Deleting Nodes

At any time you can delete a node. Note that:

- Instead of deleting a node, you may need to unmanage it. For more details, refer to the section [Managing or Unmanaging Nodes](#).
- You cannot delete a deployed node. In the column *Application name*, a deployed node belongs to a  deployed application.

You must either delete the node itself, delete the pool it belongs to, or dissociate the relevant GSLB server from the parent application. Once the GSLB server is dissociated from the application, the dedicated node deployment line is no longer listed. For more details, refer to the section [Editing Applications](#).

To delete a node

1. In the sidebar, go to  **Application > Nodes**. The page **All nodes** opens.
2. Filter the list if need be.
3. Tick the node(s) you want to delete.
4. In the menu, click on . The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The node is no longer listed, the deployed node lines are deleted as well.

Part X. Guardian

Guardian offers adaptive security to DNS cache and recursive services by detecting threats and activating adapted counter measures to ensure DNS services continuity and attack mitigation.

Guardian operates in a secured framework, with the cache separated from the recursive DNS engines. It performs a continuous real-time analysis of the inbound and outbound traffic and therefore offers complete DNS Transactions Inspection (DTI).

This analysis is especially reliable if a large number of queries is cached. By default, Guardian caches all the answers to clients' queries. The answers not cached yet, are sent to the DNS resolver and transmitted to Guardian, that caches them and sends them back to the client. That way, Guardian has an ever-growing number of queries to refine its analysis of the network and avoid potential threats.

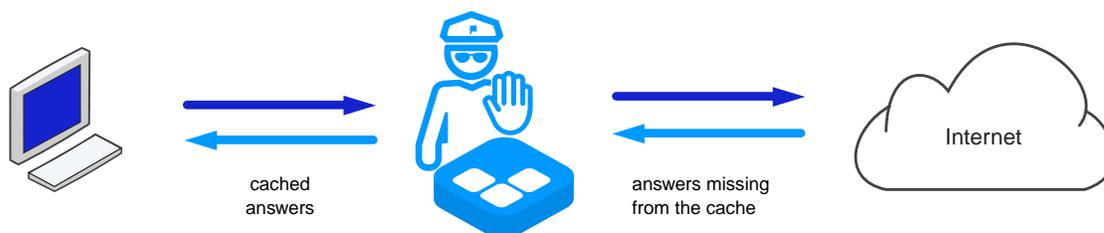


Figure 202. Guardian default behavior

Guardian is managed from the graphical user interface and via command-line interface:

1. From the **GUI**, you can enable and disable the service, configure its listening interface(s), set Guardian parameters, manually flush clients and cache, manage and configure policies and triggers and monitor dedicated charts and analytics.
2. Via **CLI**, you can connect to Guardian using SSH to display its configuration, monitor its cache and client statistics.

This part details the steps to configure Guardian and manage the functionalities it offers:

- [Configuring DNS Guardian](#): comply with the prerequisites and take the limitations into account before configuring the service.
 - [Managing Guardian Configuration](#): set Guardian parameters so that it meets your needs.
 - [Managing Guardian Cache](#): once Guardian is configured, manage its cache.
 - [Managing Guardian Statistics](#): monitor Guardian server and client statistics.
 - [Managing Guardian Protection](#): manage the actions to be triggered to protect Guardian server.
-

Chapter 54. Configuring Guardian

To configure the module Guardian and use it to its fullest, you must:

- Meet the [prerequisites](#).
- Take into account the [limitations](#).
- Configure and enabling the [service DNS Guardian](#).

Prerequisites

- A SOLIDserver appliance fully dedicated to DNS resolution in version 7.1 or higher.
 - An appliance with at least 8 GB of RAM:
 - Either a Blast-series hardware appliance, or any hardware model, except SDS-260 and SDS-3300, with a license including the option Guardian.
 - Or a virtual appliance with a license including the option Guardian. It must be configured with an *intel* network card (*em**, *ig**, *igb**, *ix**, *ixg**, *ixv** or *ixl**).
- SDS-4000 Blast virtual appliances require a PCI Pass-through to a physical Intel X520 NIC in virtual environments¹.
- A license key on every appliance where DNS Guardian is enabled.
 - An appliance properly configured:
 - The communication between the server and its clients has to be over UDP and/or TCP.
 - Both DNS and Guardian dedicated services must be running to ensure Guardian is used at its full potential. Make sure that either the services *DNS Guardian* and *DNS server (named)* or the services *DNS Guardian* and *DNS server (unbound)* are both enabled and started.
 - The appropriate rights and resources configured for end users.
 - All objects in the module *Guardian* are visible to any user by default but end-users cannot manage them if they do not belong to a group granted access to the policies and triggers *Guardian* rights.
 - To be able deploy a policy on a Guardian server, users must belong to a group that is granted the right *Display: DNS servers list* and has the server among its resources.
 - To manage Guardian parameters in the module *DNS*, users must belong to a group that is also granted the rights *Add: Guardian parameters*.

Note that Guardian is designed to express its full potential when combined with SDS-Blast series hardware appliances, whose recursive service can manage up to 17 million QPS.

Limitations

- Guardian servers in version lower than 7.1 can only be configured and managed through CLI. If the server is in a lower version, you will not be able to manage it from the GUI.
- DNS Guardian does not allow managing more than 8 views.
- DNS Guardian does not support VLAN tagging on the network interfaces it uses.

¹As in VMware environments: <https://kb.vmware.com/s/article/1010789>

- DNS Guardian does not allow capturing the network traffic answered by the service.
- On servers both recursive and authoritative, DNS Guardian does not answer queries on authoritative records, the resolver answers. For more details, refer to the section [Configuring DNS Guardian on Authoritative and Recursive Server](#).
- The module *Dashboards* does not include a Guardian dedicated page.

Enabling the Service DNS Guardian

To enable DNS Guardian properly you need to:

1. Configure the listening interfaces, as detailed in the procedure below.
2. Enable the service DNS Guardian.
3. Apply the configuration.

Note that if your license includes both DNS Guardian and DNS GSLB, you must configure the line *DNS Guardian / GSLB server* as both features rely on the same service.

To configure listening interfaces and enable DNS Guardian

Only users of the group *admin* can perform this operation.

1. Configure DNS Guardian listening interface

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **System**, click on **Services configuration**. The page **Services configuration** opens.
- c. In the column **Name**, click on **DNS Guardian** or **DNS Guardian / GSLB server**. The wizard **DNS Guardian configuration** or **DNS Guardian & GSLB server configuration** opens.
- d. In the list **Available interfaces**, select the interface of your choice and click on . The interface is moved to the list **Selected interfaces**.

Each interface is listed *<interface-name> (<MAC-address>)*, whether it is active or not. Only *Intel* network interfaces are listed as no other interface card can be configured for the service.

- e. Repeat this action for as many interfaces as you need.

To remove an interface from the list **Selected interfaces**, select it and click on . The interface is moved back to the list **Available interfaces**.

- f. Click on to complete the operation. The report opens and closes.

2. Enable the service DNS Guardian

- a. In the column **Name**, look for **DNS Guardian** or **DNS Guardian / GSLB server**.
- b. In the column **Enabled**, click on the link **Disabled** to enable the service. The wizard opens.
- c. Click on to complete the operation. The wizard closes. The page is visible again.

3. Apply your configuration

- a. Right now your configuration is pending. In the menu, select  **Tools > Apply configuration** to save your changes. The wizard **Commit the system configuration changes** opens.
- b. Click on to complete the operation. The page refreshes.

If the service is marked  *Warning*, it might mean that the system has not enough memory installed, that a configuration file is corrupt or that the license is not valid.

Once the service is enabled, on the page *All servers*, your local EfficientIP DNS server is marked  *Enabled* in the column **Guardian/GSLB**.

When the configuration is set, all the cached queries are saved in the file `/data1/dns-blast_cache.dump` which is used to answer clients queries. This file matches the server configuration and contains:

- Each cached query.
- All the ACLs configured on the server; they are listed as a set of IP addresses/network addresses that are granted or denied access.
- For BIND servers, the views match-clients and match-destinations configurations are saved as well.

In addition, in the GUI:

- *DNS Guardian* statistics panels are available on the DNS server properties page. For more details, refer to the section [Monitoring Guardian Statistics from the GUI](#).
- The analytics *DNSTOP* are no longer available on the page *Analytics*, instead *Guardian* data is available. For more details, refer to the section [Monitoring Guardian Analytics](#).

You can fully disable DNS Guardian from the GUI like any other service. Keep in mind that when you disable the service:

- DNS Guardian cache is saved automatically, as disabling a service also stops it.
- Enabling the service again will start it as well, which will restore the latest version of the cache.
- In the GUI:
 - The *Guardian* statistics panels on the properties page are no longer displayed.
 - The analytics *Guardian* are no longer available on the page *Analytics*, instead the *DNSTOP* data is available.

Note that if you stop the service, but do not disable it, all the analytics and statistics dedicated panels and list are still available but the information is outdated and therefore no longer reliable.

You can disable the Guardian protection to keep filling the cache without arming any trigger or switching to Rescue mode. For more details, refer to the section [Enabling Guardian Protection](#).

To stop or disable the service *DNS Guardian*, refer to the section [Enabling or Disabling a Service](#) in the chapter *Configuring the services*.

Chapter 55. Managing Guardian Configuration

Once you configured the service as detailed in the chapter [Configuring DNS Guardian](#), you can manage its configuration.

From the GUI, you can edit, display and save Guardian configuration. By default, the Rescue mode is enabled and configured. To set your own values, refer to the section [Configuring Guardian Rescue Mode](#) in the chapter *Managing Guardian Protection*.

Browsing Guardian Configuration

Since version 7.1, you can display Guardian configuration from the GUI.

To display the Guardian parameters you configured in the GUI

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on . The properties page opens.

In the panel **Options**, the section **Advanced options** contains the parameters with values different from the default values.

The status of the parameter is displayed.

For more details, refer to the sections [Understanding Guardian Parameters](#) and [Understanding Guardian parameter Statuses](#).

Understanding Guardian Parameters

From the GUI, you can display all the parameters configured for Guardian.

The parameters are described in the following table:

Table 55.1. Guardian configuration parameters

Configuration parameter	Description
Additional time keep in cache = <sec>	An additional period of time, in seconds, added to the cached records before letting a record expire.
Answerlog = <0 1 2>	The answerlog activation status, either disabled (0), enabled (1) or enabled only for client log (2). It is disabled by default. For more details, refer to the section Enabling Querylog and Answerlog on Triggers .
Blast=<0 1>	DNS Guardian cache activation status, either enabled (1) or disabled (0). Depending on your license, it can be enabled by default. For more details, refer to the section Enabling Guardian Protection .
Clean cache<0-9>.cmd=<filter>	The command to clean DNS Guardian cache according to the filter you set. These ten cache cleaning commands are automatically executed in order, when the automatic clearing of the cache is launched. For more details, refer to the section Clearing Guardian Cache Automatically .

Managing Guardian Configuration

Configuration parameter	Description
Clean client<0-9>.cmd=<filter>	The command to clean DNS Guardian clients data according to the filter you set. These ten client cleaning commands are automatically executed in order, when the automatic clearing of the clients is launched. For more details, refer to the section Clearing Guardian Client Statistics Automatically .
Client stats=<0 1>	The client statistics retrieval status, either disabled (0) or enabled (1). It is enabled by default. These statistics are based on the IP address of the clients, for more details refer to the section Managing Guardian Client Statistics .
List<0-7>.name=<list-name>	The configuration parameters of Guardian restriction lists. For more details refer to the section Managing Guardian Lists .
List<0-7>.request xfer=<dig-parameters>	
List<0-7>.save=<0 1>	
List<0-7>.zone name=<zone-name>	
List log=<0 1 2>	The list log activation status, either disabled (0), enabled (1) or enabled only for client log (2). It is disabled by default. For more details, refer to the section Enabling Querylog, Answerlog and List Log on Filtered Views .
List redirect aaaa=<ipv6-address>	The redirection IPv6 address for clients queries associated with the view action <i>Redirect</i> . For more details, refer to the sections Filtering Guardian Views Using Lists .
List redirect a=<ipv4-address>	The redirection IPv4 address for clients queries associated with the view action <i>Redirect</i> . For more details, refer to the section Filtering Guardian Views Using Lists .
List redirect ttl=<seconds>	The TTL of <i>List redirect aaaa</i> and <i>List redirect a</i> . For more details, refer to the section Filtering Guardian Views Using Lists .
Max cache entries=<number>	The maximum number of cache entries stored. Depending on the appliance memory, setting a large number might prevent it from running smoothly. For more details, refer to the section Clearing Guardian Cache Automatically .
Max clean cache percent=<percent>	The percentage of <i>Max cache entries</i> that defines when the cleaning cache commands <i>Clean cache<0-9>.cmd</i> are launched. Its default value is 95%. For more details, refer to the section Clearing Guardian Cache Automatically .
Max clean client percent=<percent>	The percentage of <i>Max client entries</i> that defines when the cleaning cache commands <i>Clean client<0-9>.cmd</i> are launched. Its default value is 95%. For more details, refer to the section Clearing Guardian Cache Automatically .
Max client entries=<number>	The maximum number of client entries stored. Depending on the appliance memory, setting a large number might prevent it from running smoothly. For more details, refer to the section Clearing Guardian Client Statistics Automatically .
Min clean cache percent=<percent>	The percentage of <i>Max cache entries</i> that you want to keep in the cache. Keep in mind that the more cached data you have, the more efficient your responses are and the more reliable your statistics and analytics are. For more details, refer to the section Clearing Guardian Cache Automatically .
Min clean client percent=<percent>	The percentage of <i>Max client entries</i> that you want to keep in the cache. Keep in mind that the more cached data you have, the more efficient your responses are and the more reliable your statistics and analytics are. For more details, refer to the section Clearing Guardian Client Statistics Automatically .
Min ttl cache=<seconds>	A way to enforce a minimum TTL for all the records that Guardian receives from the local recursive server before sending them to the DNS clients. The minimum only applies to records that have

Managing Guardian Configuration

Configuration parameter	Description
	a TTL lower than the one you set for the parameter. By default, the parameter is disabled (0). The maximum value accepted is 4 294 967 295 seconds.
Quarantine redirect aaaa=<ipv6-address>	The redirection IPv4 address for clients queries associated with the trigger action <i>Quarantine redirect</i> . For more details, refer to the section Managing Triggers .
Quarantine redirect a=<ipv4-address>	The redirection IPv6 address for clients queries associated with the trigger action <i>Quarantine redirect</i> . For more details, refer to the section Managing Triggers .
Quarantine redirect ttl=<seconds>	The TTL of both the <i>Quarantine redirect aaaa</i> and <i>Quarantine redirect a</i> records.
Querylog=<0 1 2>	The querylog activation status, either disabled (0), enabled (1) or enabled only for client log (2). It is disabled by default. For more details, refer to the section Enabling Querylog and Answerlog on Triggers .
Recursive=<0 1 2>	Guardian server recursion configuration, disabled (0), recursive (1) or both recursive and authoritative (2). It is enabled (1) by default. For more details, refer to the section Configuring DNS Guardian on Authoritative and Recursive Server . Setting this parameter to 0 is not recommended.
Rescue detection=<0 1 2>	The Guardian rescue mode detection configuration, either disabled (0), enabled (1) or forced (2). It is enabled (1) by default. For more details, refer to the section Managing DNS Guardian Rescue Mode .
Rescue high rec packet=<number>	The configuration parameters of Guardian Rescue mode. For more details refer to the section Managing DNS Guardian Rescue Mode .
Rescue min rec packet=<number>	
Rescue ratio servfail=<number>	
Rescue servfail qps=<number>	
Rescue time=<seconds>	
Rescue ttl=<seconds>	
Rescue unanswered rate=<number>	
Servfail diff=<0 1>	A way to identify and differentiate cached records based on the rcode they return the first time they are queried. If enabled, this parameter force DNS Guardian to answer clients querying cached records even if the local recursive server sent out a <i>SERVFAIL</i> . In this case, the records are returned with a TTL 30 seconds. This parameter only applies to records that previously sent out any other rcode than a <i>SERVFAIL</i> , and is enabled (1) by default. You can disable this parameter, for more details refer to the section Ignoring the SERVFAIL Error Message Differences .
Shared cache key=<key>	The cache sharing configuration parameters. For more details, refer to the section Sharing the Cache of Several Guardian servers .
Shared cache mcast addr=<ip-address>	
Shared cache mcast port=<port-number>	
Shared cache mcast source addr=<ip-address>	
Shared cache mcast source port=<port-number>	
Shared cache mcast ttl=<0 1>	
Shared cache=<0 1>	
Tcp passthru clients=<ip-address>	If the TCP protocol is used, a way to define the list of the IP address of appliances that Guardian will not have any impact on, separated by semi-colons. For instance, it can be useful to prevent the DNS

Configuration parameter	Description
	Guardian to have an impact on secondary DNS servers and managers.
View<0-7>.list id filter=<any all none default> <space-separated-list-ids> <action>	The configuration parameters of DNS Guardian views. For more details, refer to the section Managing Guardian Views .
View<0-7>.nat destination=<none network-address/prefix>	

Understanding Guardian parameter Statuses

The status displayed for the parameters in the section **Advanced options** of the panel **Options** provides information regarding the parameters with a value different from the default one.

Table 55.2. Guardian Parameter Statuses

Status	Description
<No status>	The parameter has already been pushed onto the server.
<i>From smart</i>	The parameter value was set at the smart level and is different from the default value.
 <i>Delayed create</i>	The parameter is being pushed onto the server.

Editing Guardian Configuration

You can edit Guardian configuration at any time.

Note that you can define Guardian parameters on a smart architecture so that the servers in the smart inherit them. At server level, you can override the inherited values.

If you set a parameter via CLI, after a couple of minutes, it appears in the GUI.

To edit Guardian configuration

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on . The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard. The list of all Guardian parameters is displayed.
5. In the drop-down list **Display options(s)**, select which parameters to display: *all, using non-default values, using default values* or *different from smart*.
6. For the parameters of your choice, select or type in a value. For more details, refer to the table [DNS Guardian configuration parameters](#).

Note that if you change the filter in the drop-down list **Display option(s)**, and if you have changed the value of a parameter that is not displayed anymore, it is not saved.

7. Click on **OK**. The report works for a while and closes.

In the panel **Options**, the parameters with a value different from the default one are displayed.

Configuring Guardian on Recursive and Authoritative Server

You can enable Guardian on a DNS server that is recursive and authoritative. Note that:

- The RRL (Response Rate Limiting) is still valid with Guardian.
- We recommend avoiding to leave Guardian authoritative on public network. The feature is designed to be used on private networks.

In order to always give an accurate answer, the authoritative queries are always sent to the resolver even if an answer to this query was cached previously. Therefore, an authoritative query always results in a cache miss, which has an impact on the analytics, the statistics and the triggers detection.

To lower this impact on the triggers detection you can create a list containing the domain names for which the server is authoritative and update the trigger for the *cache_miss*. For more details, refer to the sections [Adding Triggers Relying on Lists Metrics](#) and [Editing Triggers](#).

Note that you have a smart architecture that manages other authoritative servers and a Guardian server, it is recommended to configure the Guardian server to be used on recursive an authoritative server.

To activate Guardian on recursive and authoritative server

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⋮**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. In the field **Recursive**, select *both (2)*.
7. Click on **OK** to complete the operation. The wizard closes.

Chapter 56. Managing Guardian Cache

Once you configured the service as detailed in the chapter [Configuring DNS Guardian](#), you can manage its cache:

- Via SSH, you can connect to Guardian to display the cache content, reset it, save it, restore it, clear it and share it with other Guardian servers - incrementally or as a whole.
- From the GUI, you can clear Guardian cache.

The cache data is then used at the server and client levels to generate statistics that you can manage via CLI and visualize from the GUI. For more details, refer to the chapter [Managing Guardian Statistics](#).

Displaying Guardian Cache Content

You can display, filter and sort as many entries in the cache as needed. Below is an example of the command result limited to the 10 most queried entries:

```
DNS Blast> show cache limit=10
```

Used	Query	RCODE	Type	TTL	EOL	Last-used
213201	license.intranet	NOERROR	A	3600	1167	0
35	20	L	-			
180971	apidata.googleusercontent.com	NOERROR	A	267	57	210
281	1438	L	-			
180088	apidata.googleusercontent.com	NOERROR	AAAA	300	-827	1125
145	1487	L	-			
80249	vbox.validation.intranet	NOERROR	A	600	143	457
266	160	L	-			
67366	debian1.intranet	NOERROR	A	3600	1684	24
46	0	L	-			
57276	core.intranet	NOERROR	A	600	56	0
265	180	L	-			
54283	clients4.google.com	NOERROR	A	176	169	5
982	14704	L	-			
52764	clients6.google.com	NOERROR	A	144	139	3
988	10563	L	-			
49970	mail.google.com	NOERROR	A	299	61	31
560	7561	L	-			
47958	play.google.com	NOERROR	A	300	32	14
527	16102	L	-			

```
RCODE RETURNED:
```

RCODE	Total	Percent
NOERROR	62362	(81.61%)
SERVFAIL	1481	(1.94%)
NXDOMAIN	12576	(16.46%)

```
RR TYPE:
```

Type	Total	Percent
A	67047	(87.74%)
NS	34	(0.04%)
CNAME	10	(0.01%)
SOA	76	(0.10%)
PTR	1815	(2.38%)
MX	26	(0.03%)
TXT	555	(0.73%)
AAAA	6739	(8.82%)
SRV	80	(0.10%)
NAPTR	3	(0.00%)
DS	2	(0.00%)
SSHFP	12	(0.02%)
DNSKEY	3	(0.00%)
ANY	16	(0.02%)

```
ORIGIN:
```

```

Origin |      Total | Percent
  L   |    76419  | (100.00%)

76419 entries found (5115k total utilization  avg_pkt_size: 91)

search time: 20ms

```

The results of the command are described in the following table:

Table 56.1. Guardian cache statistics

Data returned	Description
Used	The number of times the query was made, or used. The results are listed from most queried to least queried.
Query	The name of the queried domain.
RCODE	The query response status, it indicates possible errors.
Type	The type of records queried.
TTL	The record time to live in seconds.
EOL	The record expiration time in seconds, or end of life. A negative <i>eol</i> means that the record has expired.
Last-used	The period in seconds since the record was queried for the last time.
R-time	The cumulated time in recursion in milliseconds.
AVG R-time	The average time in recursion in milliseconds.
C-miss	The number of queries for which the TTL had expired.
View	The view the record belongs to. If there are no views on the server, the view <i>0</i> is the only one displayed.
Origin	The record origin, either L for Local or S for Sent. For more details, refer to the sections Sharing the Cache of Several Guardian servers and Sending the Cache Content to Another Guardian server .
OPT	The option set in the query, None, DNSSEC or EDNS.
RCODE RETURNED	The number and percentage of RCODEs returned, by type: <i>NOERROR</i> , <i>SERVFAIL</i> , <i>NXDOMAIN</i> ..
RR TYPE	The number and percentage of resource records returned, by type: <i>A</i> , <i>AAAA</i> , <i>NS</i> , <i>CNAME</i> ...
ORIGIN	The number and percentage of queries, by origin: either L for Local or S for Sent. For more details, refer to the sections Sharing the Cache of Several Guardian servers and Sending the Cache Content to Another Guardian server .
<n> entries found	The total number of different results matching the command filter(s).
<n> total utilization	The total number of queries received matching the command filter(s).
avg_pkt_size: <number-of-bytes>	The average size of the packets returned matching the command filter(s).

To display Guardian cache statistics

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To display Guardian cache statistics, use the following command:

```
show cache
```

4. To limit the entries to be displayed using the following command:

```
show cache limit=<number>
```

- To sort the entries by the type of data returned, described in the table [Guardian cache statistics](#), type in the following command. By default, the command displays 50 entries, in particular, the 50 most queried domains.

```
show cache order=<data-type>
```

- To display specific cache entries, use the following command with one or several of the filters from the table below. You can combine the order and limit parameters with one or several filters in one command.

```
show cache <filter>
```

Table 56.2. Available filters for Guardian cache statistics

Filter	Description
view=<0-7>	Filter the results using views. As you cannot have more than 8 views with Guardian, the views number is an integer between 0 and 7 that corresponds to the order you set for the local recursive server views. For more details, refer to the section Editing the Order of the Views . If you did not configure any view, all the results returned belong to view=0.
expired=<0 1>	Filter the results using records TTL expiration status. <code>expired=0</code> returns the cached records with a valid TTL, <code>expired=1</code> displays the expired records in the cache. If you do not specify this filter, all the records are returned no matter their TTL.
eo!(> <)<number-of-seconds>	Filter the results using the record expiration time. Use > or < to narrow down your search. To exclude expired values, use the <code>expired</code> filter.
<error(= !=)ERROR>	Filter the results using the errors returned. Use the format <code>error=<ERRORTYPE></code> or <code>error!=<ERRORTYPE></code> to narrow down your search. Accepted error types are: NOERROR, FORMERR, SERVFAIL, NXDOMAIN, NOTIMP or REFUSED.
type=<RRTYPE>	Filter the results using the record types. All record types are accepted.
<origin=L S>	Filter the results using the record origin: cached by the local Guardian server, <code>origin=L</code> , or shared/sent by another Guardian server, <code>origin=S</code> .
opt=<- DNSSEC EDNS>	Filter the results using the option the records are configured with. The accepted options are: - (no option was set), DNSSEC and EDNS.
used(> < = !=)<number>	Filter the results using the number of times the records were queried, i.e. the number of hits. Use >, <, = or != to narrow down your search.
lastused(> < = !=)<number-of-seconds>	Filter the results using the last time the entries were queried in seconds. Use >, <, = or != to narrow down your search.
ttl(> < = !=)<number-of-seconds>	Filter the results using the record TTL value, in seconds. Use >, <, = or != to narrow down your search.
<regex filter>	Filter the results using regular expressions.
<domain>	Filter the results using a full or partial domain name. For instance, the command <code>show cache view=0 mit.edu</code> only returns hits for queries of the domain <code>mit.edu</code> within the first - or only - view of the local recursive server. The domain name is always the last parameter of the command. You can also exclude the domain name of your choice following the syntax <code>!<domain></code> .

Resetting Guardian Cache

You can reset the number of hits for some or all of the listed *used* entries in the cache.

To reset Guardian cache entries

- Open a shell session on your appliance using *root* credentials.

2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To completely reset the current version of Guardian cache, use the following command:

```
reset cache
```

4. To reset only a part of Guardian cache, use the filters detailed in the table [Available filters for Guardian cache statistics](#) and the following command:

```
reset cache <filter>
```

Saving Guardian Cache

You can save Guardian cache at any time. The last cache version can be used as a point of restoration.

Keep in mind that every time you stop or restart DNS Guardian, the cache is automatically saved or restored.

To save Guardian cache

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To save Guardian cache, use the following command:

```
save cache
```

The cache is saved in the file `/data1/dnsblast_cache.dump`. Note that you cannot browse the content of this file, it can only be used for restoration purposes.

Restoring Guardian Cache

Whether the cache was saved automatically or at a specific time you can use a command to restore the latest saved version.

The file `/data1/dnsblast_cache.dump` contains the latest version of the cache and is automatically used when you restore the Guardian server cache.

To restore Guardian cache

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To restore the last saved version of Guardian cache, use the following command:

```
load cache
```

Forcing Cache Entries as Expired

You can set all or part of the cache entries as expired. This operation is useful you do not want to keep using the entries in the cache but want to keep the statistics related to them.

To set Guardian cache entries as expired

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To set Guardian cache entries as expired, use the following command:

```
expire cache
```

4. To set only a part of Guardian cache as expired, use the filters detailed in the table [Available filters for Guardian cache statistics](#) and the following command:

```
expire cache <filter>
```

Clearing Guardian Cache Manually

At any time, you can clear records from your Guardian cache, that is to say delete some or all the entries it contains.

Clearing Guardian Cache Manually via CLI

A command indicates the number of entries that were deleted. Below is an example using a regex filter.

```
DNS Blast> clear cache (mit.edu|ietf)
32 entries flushed
done
```

Note that you can also set parameters to automatically clear the cache to keep Guardian performances at best. For more details, refer to the section [Clearing Guardian Cache Automatically](#).

To clear Guardian cache via CLI

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To clear the cache records content, entirely or partly, use the following command:

```
clear cache
```

4. To clear only a part of the cache records content, use the filters detailed in the table [Available filters for Guardian cache statistics](#) and the following command:

```
clear cache <filter>
```

Clearing Guardian Cache Manually from the GUI

Since version 7.1, you can clear Guardian cache from the module DNS.

To clear Guardian cache from the GUI

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Select the Guardian server(s) of your choice.
3. In the menu, select **Edit > Command > Flush cache**. The wizard **Flush DNS Cache** opens.
4. Tick the box **Flush Guardian cache**. The field **Cache entries regex** appears.
5. In the field **Cache entries regex**, specify the cache entries you want to flush via a regular expression filter.

Keep in mind that to use special characters like the dot "." as a literal in a regex, it must be escaped with a backslash "\".

6. Click on **OK** to complete the operation. The wizard closes.

Clearing Guardian Cache Automatically

You can use a set of parameters in Guardian configuration to automate the cache entries clearing process.

Three parameters allow to define a maximum number of entries in the cache, the threshold to be reached to start the automatic clearing of the cache and the maximum number of remaining entries after the clearing. When the maximum percentage of entries is reached, the clearing commands are automatically executed in ascending order until the number of remaining queries equals or is lower than the minimum percentage of entries. You can configure up to 10 clearing commands matching the filters of your choice.

The automated clearing is configured using the parameters below:

- **Max cache entries** sets the maximum of entries that you want to store in the cache. Its default value depends on the memory of each appliance:

Table 56.3. Max cache entries default values

SOLIDserver model	Max cache entries default value
SDS-550	500000
SDS-1100	500000
SDS-2200	500000
SDS-3300 & Blast series	1000000

- **Max clean cache percent** sets the percentage of *Max cache entries* that is to be reached to launch the execution of the clearing commands *Clean cache<0-9>.cmd*.
- **Min clean cache percent** sets the percentage of *Max cache entries* to keep in the cache after executing the clearing commands *Clean cache<0-9>.cmd*.

Once *Max clean cache percent* is reached, the cleaning command *Clean cache0.cmd* is executed. If the cache still contains more entries than the value set with *Min clean cache percent*, the command *Clean cache1.cmd* is executed. All the commands *Clean cache<0-9>.cmd* are

executed in order until the percentage of entries remaining in the cache is equal or lower than the value set with *Min clean cache percent*.

- **Clean cache<0-9>.cmd** parameters set 10 different clearing commands that respect the clearing filters of your choice.

The clearing commands are only executed if:

- The parameter *Clean cache<0-9>.cmd* has a value.
- The number of entries left in the cache is higher than the value set with *Min clean cache percent*.

By default, only the first three clearing commands are configured.

Note that you can also clear some or all the cache entries manually. For more details, refer to the section [Clearing Guardian Cache Manually](#).

To automate Guardian cache clearing

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To configure the clearing commands, in the fields **Clean cache<0-9>.cmd**, use the filters detailed in the table [Available filters for Guardian cache statistics](#).
7. In the field **Max cache entries**, specify the maximum number of entries in Guardian cache.
8. In the field **Max clean cache percent**, specify the percentage of the maximum number of cache entries that will start the automatic clearing of the cache.
9. In the field **Min clean cache percent**, specify the percentage of the maximum number of cache entries that you want to keep in Guardian cache after the clearing.
10. Click on **OK** to complete the operation. The wizard closes.

Sharing the Cache between Several Guardian servers

If you purchased several DNS Guardian licenses, Guardian servers can share their caches:

- Your servers can automatically share their latest cache entries, as detailed in the section [Setting Up Incremental Cache Sharing Between Several Guardian servers](#).
- You can send the entire cache content from one Guardian server to another, as detailed in the section [Sending the Cache Content to Another Guardian server](#).

Setting Up Incremental Cache Sharing Between Several Guardian servers

If you set up an incremental cache sharing, any new record in the cache is automatically added to the cache of the other servers configured.

Once the sharing is set on all servers:

1. They can automatically update their cache content using a common multicast configuration.
2. You can send the entire content of your cache to a specific Guardian server using a unicast IP address. For more details, refer to the section [Sending the Cache Content to Another Guardian server](#).

To set the cache sharing between Guardian servers, on each of them you must:

- Specify a common multicast IP address.
- Specify a common multicast port.
- Specify a common sharing key.
- Enable the sharing service.
- Set the TTL of the IP packets sent. By default, it is set to 1.

To set up an incremental cache update between Guardian servers via CLI

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. Set the multicast IP address for the sharing using the following command:

```
set shared_cache_mcast_addr=<multicast-IP-address>
```

4. Set the multicast port used for the cache sharing using the following command:

```
set shared_cache_mcast_port=<multicast-port>
```

5. Set the sharing key using the following command:

```
set shared_cache_key=<key>
```

6. Enable the sharing using the following command:

```
set shared_cache=1
```

7. Set the shared IP packets TTL using the following command:

```
shared_cache_mcast_ttl=<TTL>  
  
// Its default value is 1
```

8. Repeat the steps 3 to 7 using the exact same IP address, port and key on all the Guardian servers that should send and receive the new entries in their cache.

To set up an incremental cache update between several Guardian servers from the GUI

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.

6. For the parameter **Shared cache mcast addr**, specify the multicast IP address for the sharing.
7. For the parameter **Shared cache mcast port**, specify the multicast port used for the cache.
8. For the parameter **Shared cache key**, specify the sharing key.
9. To enable the sharing, for the parameter **Shared cache**, select *yes (1)*.
10. For the parameter **Shared cache mcast ttl**, specify the shared IP packets TTL. Its default value is *1*.
11. Click on to complete the operation. The wizard closes.
12. Repeat these steps using the exact same IP address, port and key for all the Guardian servers that should send and receive the new entries in their cache.

The entries on the local Guardian are marked *origin: L* while they are marked *origin: S* on the receiving one. For more details, refer to the section [Displaying Guardian Cache Content](#).

Sending the Cache Content to Another Guardian server

If you enabled Guardian on several appliances, once your cache has all the entries you need, you can send its entire content to another Guardian server. Using a dedicated command, you can send all the cache entries from the source server to a receiving server, which then adds the missing entries to its cache and/or updates the existing ones.

To send the content of your cache you must:

- Configure the cache sharing parameters: the multicast IP address, port and sharing key. It is impossible to send the content of the cache from one Guardian server to the other without these parameters. Once properly configured, the sharing must be enabled (*Shared cache* is set to *yes (1)*). For more details, refer to the section [Sharing the Cache of Several Guardian servers](#).
- Use the IP address of the target Guardian server. This command, unlike cache sharing, is a one-time operation that uses a unicast IP address. Any IP address configured on the target server can be used.

To share the cache with another Guardian server

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```
3. Make sure that cache sharing is configured and enabled. For more details, refer to the section [Sharing the Cache of Several Guardian servers](#).
4. To send the entire content of your cache, use the following command:

```
share cache <target-DNS-Guardian-server-IP-address>
```
5. If you do not want to incrementally update the cache of both servers with the new records, you need to disable the cache sharing parameters once the target Guardian has received the cache entries.

The entries are marked *origin: L* on the local DNS Guardian and *origin: S* on the receiving one.

Chapter 57. Managing Guardian Statistics

Once you configured the service as detailed in the chapter [Configuring DNS Guardian](#), you can manage and monitor Guardian statistics.

DNS Guardian statistics gather all the metrics related to the client-server interactions and can be fully managed via CLI. SOLIDserver also provides a selection of graphs as well as detailed analytics to help you monitor suspicious traffic from the GUI. For more details, refer to the section [Monitoring Guardian from the GUI](#).

Note that you can also monitor Guardian data yourself via SNMP. For more details, refer to the section [Monitoring Using SNMP](#).

Managing Guardian server Statistics

Guardian server statistics allow you to have a detailed overview of the traffic processed by the server. They are based on the analysis of Guardian cache and clients data to display the type, number, size and recursion time of all queries and answers. For more details, refer to the chapter [Managing Guardian Cache](#) and the section [Managing Guardian Client Statistics](#).

Server statistics can be reset at any time. For more details, refer to the section [Resetting Guardian server Statistics](#).

Displaying Guardian server Statistics

At any time you can display your Guardian server statistics as in the example below:

```
DNS Blast> show stats

 2010M Cache hit
17311k Cache missedmak
17311k Cache missed but exists (expired)
   0 Cache missed, never queried
   0 Cache hit in quarantine mode
   0 Cache missed in quarantine mode
   0 Cache missed but exists (expired) in quarantine mode
   0 Cache missed, never queried in quarantine mode
   0 Cache missed in quarantine_redirect mode
   0 Cache hit in rescue mode
   0 Cache missed in rescue mode
   0 Cache missed but exists (expired) in rescue mode
   0 Cache missed, never queried in rescue mode
   0 Blocked queries
   0 Ratelimited queries
63021 Cache size
   0 Client size
   0 Client added
   0 Client deleted
   0 Server size
2027M Sent DNS answers
172GB Sent DNS answers (bytes)
2027M Received DNS queries
141GB Received DNS queries (bytes)
17311k Queries sent to the local recursive server
1172MB Queries sent to the local recursive server (bytes)
17072k Answers sent by the local recursive server
1405MB Answers sent by the local recursive server (bytes)
   0 Sent DNS answers in quarantine mode
   0B Sent DNS answers in quarantine mode (bytes)
   0 Sent DNS answers in quarantine_redirect mode
```

```

0B Sent DNS answers in quarantine_redirect mode (bytes)
0 Sent DNS answers in rescue mode
0B Sent DNS answers in rescue mode (bytes)
0 Received Invalid DNS queries
0B Received Invalid DNS queries (bytes)
743661 Received SERVFAIL from local recursive server (rcode changed)
0 List matched queries
0 Discarded outgoing DNS packets
0 Discarded outgoing DNS packets (received by local recursive server)
0 Discarded incoming packets
0 Discarded DNS queries (sent to local recursive server)
15782k Recursive queries with latency < 10ms
1026k Recursive queries with latency >= 10ms and < 100ms
67408 Recursive queries with latency >= 100ms and < 500ms
393 Recursive queries with latency >= 500ms and < 800ms
16177 Recursive queries with latency >= 800ms and < 1600ms
180065 Recursive queries with latency >= 1600ms
2026M Sent NOERROR rcode
0 Sent FORMERR rcode
743661 Sent SERVFAIL rcode
0 Sent NXDOMAIN rcode
0 Sent NOTIMP rcode
0 Sent REFUSED rcode
0 Sent YXDOMAIN rcode
0 Sent YXRRSET rcode
0 Sent NXRRSET rcode
0 Sent NOTAUTH rcode
0 Sent NOTZONE rcode
0 Received NOERROR rcode from local recursive server
0 Received FORMERR rcode from local recursive server
0 Received SERVFAIL rcode from local recursive server
0 Received NXDOMAIN rcode from local recursive server
0 Received NOTIMP rcode from local recursive server
0 Received REFUSED rcode from local recursive server
0 Received YXDOMAIN rcode from local recursive server
0 Received YXRRSET rcode from local recursive server
0 Received NXRRSET rcode from local recursive server
0 Received NOTAUTH rcode from local recursive server
0 Received NOTZONE rcode from local recursive server

```

The results of the command are described in the following table:

Table 57.1. Guardian server statistics

Metrics	Description
Cache hit	Number of queries answered by the cache.
Cache missed	Number of queries that cannot be answered by the cache.
Cache missed but exists (expired)	Number of queries that cannot be answered by the cache because the cache entry expired.
Cache missed, never queried	Number of queries that cannot be answered by the cache because the entry queried does not exist.
Cache hit in quarantine mode	Number of queries answered by the cache to clients in quarantine mode.
Cache missed in quarantine mode	Number of that cannot be answered by the cache to clients in quarantine mode.
Cache missed but exists (expired) in quarantine mode	Number of queries that cannot be answered by the cache to clients in quarantine mode because the cache entry expired.
Cache missed, never queried in quarantine mode	Number of queries that cannot be answered by the cache to clients in quarantine mode because the entry queried does not exist.
Cache missed in quarantine_redirect mode	Number of queries that cannot be answered by the cache and are redirected to the IP address configured in the quarantine_redirect mode.
Cache hit in rescue mode	Number of queries answered by the cache in rescue mode.
Cache missed in rescue mode	Number of queries that cannot be answered by the cache in rescue mode.

Managing Guardian Statistics

Metrics	Description
Cache missed but exists (expired) in rescue mode	Number of queries that cannot be answered by the cache in rescue mode because the records expired.
Cache missed, never queried in rescue mode	Number of queries that cannot be answered by the cache in rescue mode because the entry queried does not exist.
Blocked queries	Number of queries that are blocked, once the action <i>Block</i> has been configured and triggered for specific clients.
Ratelimited queries	Number of queries that dropped because source clients reached the configured query limit rate of 100 queries per second.
Cache size	Total number of entries in the cache.
Client size	Total number of clients (unique IP addresses) for which statistics are maintained. For more details, refer to the section Managing Guardian Client Statistics .
Sent DNS answers	Number of answers sent out by Guardian.
Sent DNS answers (bytes)	Total size of the packets sent out by Guardian.
Received DNS queries	Number of queries received by Guardian.
Received DNS queries (bytes)	Total size of the packets received by Guardian.
Queries sent to the local recursive server	Number of queries not already cached by Guardian that were sent to the local DNS recursive server.
Queries sent to the local recursive server (bytes)	Total size of the packets not already cached by Guardian that were sent to the local DNS recursive server.
Answers sent by the local recursive server	Number of queries not already cached by Guardian that were answered by the local DNS recursive server.
Answers sent by the local recursive server (bytes)	Total size of the packets not already cached by Guardian that were answered by the local DNS recursive server.
Sent DNS answers in quarantine mode	Number of answers sent out to clients in quarantine mode.
Sent DNS answers in quarantine mode (bytes)	Total size of the packets sent out to clients in quarantine mode.
Sent DNS answers in quarantine_redirect mode	Number of answers redirected to the IP address configured in the quarantine_redirect mode.
Sent DNS answers in quarantine_redirect mode (bytes)	Total size of the packets redirected to the IP address configured in the quarantine_redirect mode.
Sent DNS answers in rescue mode	Number of queries received by Guardian once it switched to rescue mode.
Sent DNS answers in rescue mode (bytes)	Total size of the packets received by Guardian once it switched to rescue mode.
Received Invalid DNS queries	Number of invalid queries received by Guardian.
Received Invalid DNS queries (bytes)	Total size of the invalid queries packets received by Guardian.
Received SERVFAIL from local recursive server (rcode changed)	Number of SERVFAIL error messages received by Guardian from the local recursive server that previously returned a different rcode for the same query. To stop taking into account these results, refer to the section Ignoring the SERVFAIL Error Message Differences .
RPZ matched queries	Number of queries that matched an RPZ entry.
Discarded outgoing DNS packets	Number of DNS packets received and discarded by Guardian.
Discarded outgoing DNS packets from local recursive server	Number of DNS packets sent to the local DNS recursive server and transmitted to the client.
Discarded incoming packets	Number of incoming packets that were discarded.
Discarded DNS queries (sent to local recursive server)	Number of packets sent to the local DNS recursive server that were discarded.

Managing Guardian Statistics

Metrics	Description
Recursive queries with latency < 10ms	Number of recursive queries which answering time is lower than 10 milliseconds.
Recursive queries with latency >= 10ms and < 100ms	Number of recursive queries which answering time is equal to 10 or lower than 100 milliseconds.
Recursive queries with latency >= 100ms and < 500ms	Number of recursive queries which answering time is equal to 100 or lower than 500 milliseconds.
Recursive queries with latency >= 500ms and < 800ms	Number of recursive queries which answering time is equal to 500 or lower than 800 milliseconds.
Recursive queries with latency >= 800ms and < 1600ms	Number of recursive queries which answering time is equal to 800 or lower than 1600 milliseconds.
Recursive queries with latency >= 1600ms	Number of recursive queries which answering time is equal to or higher than 1600 milliseconds.
Sent NOERROR rcode	Number of NOERROR error messages sent by Guardian.
Sent FORMERR rcode	Number of FORMERR error messages sent by Guardian.
Sent SERVFAIL rcode	Number of SERVFAIL error messages sent by Guardian.
Sent NXDOMAIN rcode	Number of NXDOMAIN error messages sent by Guardian.
Sent NOTIMP rcode	Number of NOTIMP error messages sent by Guardian.
Sent REFUSED rcode	Number of REFUSED error messages sent by Guardian.
Sent YXDOMAIN rcode	Number of YXDOMAIN error messages sent by Guardian.
Sent YXRRSET rcode	Number of YXRRSET error messages sent by Guardian.
Sent NXRRSET rcode	Number of NXRRSET error messages sent by Guardian.
Sent NOTAUTH rcode	Number of NOTAUTH error messages sent by Guardian.
Sent NOTZONE rcode	Number of NOTZONE error messages sent by Guardian.
Received NOERROR rcode from local recursive server	Number of NOERROR error messages received by Guardian from the local recursive server.
Received FORMERR rcode from local recursive server	Number of FORMERR error messages received by Guardian from the local recursive server.
Received SERVFAIL rcode from local recursive server	Number of SERVFAIL error messages received by Guardian from the local recursive server.
Received NXDOMAIN rcode from local recursive server	Number of NXDOMAIN error messages received by Guardian from the local recursive server.
Received NOTIMP rcode from local recursive server	Number of NOTIMP error messages received by Guardian from the local recursive server.
Received REFUSED rcode from local recursive server	Number of REFUSED error messages received by Guardian from the local recursive server.
Received YXDOMAIN rcode from local recursive server	Number of YXDOMAIN error messages received by Guardian from the local recursive server.
Received YXRRSET rcode from local recursive server	Number of YXRRSET error messages received by Guardian from the local recursive server.
Received NXRRSET rcode from local recursive server	Number of NXRRSET error messages received by Guardian from the local recursive server.
Received NOTAUTH rcode from local recursive server	Number of NOTAUTH error messages received by Guardian from the local recursive server.
Received NOTZONE rcode from local recursive server	Number of NOTZONE error messages received by Guardian from the local recursive server.

Note that:

- Some of the server statistics are used to generate charts available in the GUI, for more details refer to the section [Monitoring Guardian Statistics from the GUI](#).
- You can filter the server statistics by view.
- You can also monitor the server or view statistics to display more information, refreshed every second. For more details, refer to the section [Monitoring Guardian server Statistics in Real-Time](#).

To display Guardian server statistics

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To display Guardian server statistics, use the following command:

```
show stats
```

4. To display Guardian server statistics of a specific view, use the following command:

```
show stats view=<0-7>
```

Monitoring Guardian server Statistics in Real-Time

You can monitor, in real-time, all the server statistics of Guardian or only some information. The statistics are updated every second in the column *Last sec* while the column *Cumulative* displays the values collected since the last reset. For more details, refer to the section [Resetting Guardian server Statistics](#).

Monitoring Guardian data returns information as follows:

```
DNS Blast> monitor
DNS BLAST monitoring:
-----
Cumulative (      Avg /s ) Name
-----
 99.99% (    100.00% ) Cache hit rate
  0.00% (     0.00% ) Cache miss rate
  0.00% (     0.00% ) Block query rate
15690M (   17080k ) Cache hit
    2 (         0 ) Cache missed
    2 (         0 ) Cache missed but exists (expired)
    0 (         0 ) Cache missed, never queried
    0 (         0 ) Cache hit in quarantine mode
    0 (         0 ) Cache missed in quarantine mode
    0 (         0 ) Cache missed but exists (expired) in quarantine mode
    0 (         0 ) Cache missed, never queried in quarantine mode
    0 (         0 ) Cache missed in quarantine_redirect mode
    0 (         0 ) Cache hit in rescue mode
    0 (         0 ) Cache missed in rescue mode
    0 (         0 ) Cache missed but exists (expired) in rescue mode
    0 (         0 ) Cache missed, never queried in rescue mode
    0 (         0 ) Blocked queries
    0 (         0 ) Ratelimited queries
 1003 (         0 ) Cache size
    0 (         0 ) Client size
    0 (         0 ) Client added
    0 (         0 ) Client deleted
    0 (         0 ) Server size
15690M (   17080k ) Sent DNS answers
1403GB (   1564MB ) Sent DNS answers (bytes)
```

Managing Guardian Statistics

```

15690M ( 17080k ) Received DNS queries
1169GB ( 1303MB ) Received DNS queries (bytes)
2 ( 0 ) Queries sent to the local recursive server
166B ( 0B ) Queries sent to the local recursive server (bytes)
2 ( 0 ) Answers sent by the local recursive server
284B ( 0B ) Answers sent by the local recursive server (bytes)
0 ( 0 ) Sent DNS answers in quarantine mode
0B ( 0B ) Sent DNS answers in quarantine mode (bytes)
0 ( 0 ) Sent DNS answers in quarantine_redirect mode
0B ( 0B ) Sent DNS answers in quarantine_redirect mode (bytes)
0 ( 0 ) Sent DNS answers in rescue mode
0B ( 0B ) Sent DNS answers in rescue mode (bytes)
0 ( 0 ) Received Invalid DNS queries
0B ( 0B ) Received Invalid DNS queries (bytes)
0 ( 0 ) Received SERVFAIL from local recursive server (rcode changed)
0 ( 0 ) List matched queries
0 ( 0 ) Discarded outgoing DNS packets
0 ( 0 ) Discarded outgoing DNS packets (received by local recursive server)
0 ( 0 ) Discarded incoming packets
0 ( 0 ) Discarded DNS queries (sent to local recursive server)
2 ( 0 ) Recursive queries with latency < 10ms
0 ( 0 ) Recursive queries with latency >= 10ms and < 100ms
0 ( 0 ) Recursive queries with latency >= 100ms and < 500ms
0 ( 0 ) Recursive queries with latency >= 500ms and < 800ms
0 ( 0 ) Recursive queries with latency >= 800ms and < 1600ms
0 ( 0 ) Recursive queries with latency >= 1600ms
15690M ( 17080k ) Sent NOERROR rcode
0 ( 0 ) Sent FORMERR rcode
0 ( 0 ) Sent SERVFAIL rcode
4 ( 0 ) Sent NXDOMAIN rcode
0 ( 0 ) Sent NOTIMP rcode
0 ( 0 ) Sent REFUSED rcode
0 ( 0 ) Sent YXDOMAIN rcode
0 ( 0 ) Sent YXRRSET rcode
0 ( 0 ) Sent NXRRSET rcode
0 ( 0 ) Sent NOTAUTH rcode
0 ( 0 ) Sent NOTZONE rcode
15 ( 0 ) Received NOERROR rcode from local recursive server
0 ( 0 ) Received FORMERR rcode from local recursive server
4 ( 0 ) Received SERVFAIL rcode from local recursive server
3 ( 0 ) Received NXDOMAIN rcode from local recursive server
0 ( 0 ) Received NOTIMP rcode from local recursive server
0 ( 0 ) Received REFUSED rcode from local recursive server
0 ( 0 ) Received YXDOMAIN rcode from local recursive server
0 ( 0 ) Received YXRRSET rcode from local recursive server
0 ( 0 ) Received NXRRSET rcode from local recursive server
0 ( 0 ) Received NOTAUTH rcode from local recursive server
0 ( 0 ) Received NOTZONE rcode from local recursive server

```

The result of the command starts with statistics on the current measurement rates, as described in the table below:

Table 57.2. Guardian monitoring rate statistics

Metrics	Description
Cache hit rate	Percentage of queries answered by the cache, all modes taken together (normal, quarantine and rescue mode). This line is only returned if you monitor the statistics. For more details refer to the section Monitoring Guardian server Statistics in Real-Time .
Cache miss rate	Percentage of queries that cannot be answered by the cache, all modes taken together. This line is only returned if you monitor the statistics. For more details refer to the section Monitoring Guardian server Statistics in Real-Time .
Block query rate	Percentage of queries that are blocked. This line is only returned if you monitor the statistics. For more details refer to the section Monitoring Guardian server Statistics in Real-Time .

All the others statistics are described in the table [Guardian server statistics](#).

To monitor Guardian server statistics in real-time

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To monitor all the statistics at once, use the following command:

```
monitor
```

4. To focus the monitoring on specific statistics, described in the table [Guardian server statistics](#), use the following command:

```
monitor <filter>
```

The available filters are described below, you can combine them when you execute the command.

Table 57.3. Available filters for DNS server real-time monitoring

Filter	Description
view=<view-number>	Filter the results using the statistics of a specific view
show-qps	Filter the results using the queries per second statistics, number or size in bytes.
show-rtt	Filter the results using the round-trip or recursion time statistics. These statistics are named <i>Recursive queries with latency <range></i> .
show-rcode	Filter the results using the RCODEs statistics. These statistics are named <i>Sent <VALUE> rcode</i> .

5. Hit enter to stop monitoring.

Monitoring Triggers Statistics in Real-Time

You can monitor, in real-time, all the server statistics of triggers. Monitoring Guardian data returns information as follows:

```
DNS Blast> monitor show-trigger

DNS BLAST monitoring:
-----
Cumulative (      Avg /s ) Name
-----
54001 (      121 ) Restricted IPv4 entries
2660 (       23 ) Restricted IPv6 entries
0 (        0 ) Number of time trigger #0 was armed
0 (        0 ) Number of time trigger #1 was armed
0 (        0 ) Number of time trigger #2 was armed
0 (        0 ) Number of time trigger #3 was armed
0 (        0 ) Number of time trigger #4 was armed
0 (        0 ) Number of time trigger #5 was armed
0 (        0 ) Number of time trigger #6 was armed
0 (        0 ) Number of time trigger #7 was armed
0 (        0 ) Number of time trigger #8 was armed
0 (        0 ) Number of time trigger #9 was armed
0 (        0 ) Number of time trigger #10 was armed
...
0 (        0 ) Number of time trigger #63 was armed
```

The result of the command is described in the following table:

Table 57.4. Guardian server real-time trigger statistics

Data returned	Description
Restricted IPv4 entries	Number of IPv4 addresses for which a trigger was armed.
Restricted IPv6 entries	Number of IPv6 addresses for which a trigger was armed.
Number of time trigger #0 was armed	Number of times the first trigger, #0, has been armed.
Number of time trigger #... was armed	Number of times one of the triggers, all identified with a number between #0 and #63, has been armed.
Number of time trigger #63 was armed	Number of times the last trigger, #63, has been armed.

To monitor Guardian server triggers statistics in real-time

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To monitor all the triggers statistics at once, use the following command:

```
monitor show-trigger
```

4. Hit enter to stop monitoring.

Monitoring Guardian Cache Sharing Statistics in Real-Time

You can monitor, in real-time, all the server statistics of cache data shared among several Guardian appliances. Monitoring Guardian data returns information as follows:

```
DNS Blast> monitor show-cache-sharing
DNS BLAST monitoring:
-----
Cumulative (      Avg /s ) Name
-----
      0 (          0 ) DNS entries received via shared cache
     0B (         0B ) DNS entries received via shared cache (bytes)
      0 (          0 ) DNS entries sent via shared cache
     0B (         0B ) DNS entries sent via shared cache (bytes)
```

The result of the command is described in the following table:

Table 57.5. Guardian server real-time cache-sharing statistics

Data returned	Description
DNS entries received via shared cache	Number of queries received by the local Guardian via shared cache.
DNS entries received via shared cache (bytes)	Total size of the packets received by the local Guardian via shared cache.
DNS entries sent via shared cache	Number of queries sent to another Guardian via shared cache.
DNS entries sent via shared cache (bytes)	Total size of the packets sent to another Guardian via shared vcache.

To monitor Guardian server cache sharing statistics in real-time

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To monitor the cache sharing statistics of the server, use the following command:

```
monitor show-cache-sharing
```

4. Hit enter to stop monitoring.

Resetting Guardian server Statistics

At any time you can reset the statistics of your Guardian server to make sure that the data displayed is up to date, especially if you are [Monitoring Guardian server Statistics in Real-Time](#).

To reset Guardian server statistics

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To reset the statistics results, use the following command:

```
reset stats
```

Ignoring the SERVFAIL Error Message Differences

By default, Guardian monitors the SERVFAIL message differences as a record queried several times can have any rcode at first and then a SERVFAIL message. A command allows to force Guardian to ignore this difference.

This information is listed on the line *Received SERVFAIL from local recursive server (rcode changed)* when you execute the command `show stats`. For more details, refer to the section [Displaying Guardian server Statistics](#).

To disable the detection of SERVFAIL message differences

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⌵**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To disable the detection of SERVFAIL message differences, in the field **Servfail diff**, select *no (0)*.
7. Click on **OK** to complete the operation. The wizard closes.

Note that Guardian provides you with two Guardian tops dedicated to queries and clients returning SERVFAIL error message: *Top clients receiving SERVFAIL rcodes* and *Top queries returning SERVFAIL rcodes*. For more details, refer to the section [Monitoring Guardian Analytics](#).

In addition, you can configure Guardian switch to Rescue Mode based on the number of SERVFAIL error messages received by local recursive servers. For more details, refer to the section [Configuring Guardian Rescue Mode](#).

Managing Guardian Client Statistics

Guardian client statistics allow you to have a specific view of the server usage focused on the IP address making the queries. They are based on the analysis of Guardian cache and the server statistics in general. For more details, refer to the chapter [Managing Guardian Cache](#) and the section [Managing Guardian server Statistics](#).

Making Sure Guardian Client Statistics Are Enabled

To manage Guardian client statistics, make sure that the feature is enabled. It is enabled by default. If on the Guardian server properties page, in the panel *Options*, the parameter **Client stats** is listed, it means that its value is different from the one by default and you need to enable it.

To enable Guardian client statistics

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To enable Guardian client statistics, in the field **Client stats**, select *yes (1)*.
7. Click on **OK** to complete the operation. The wizard closes.

Displaying Guardian Client Statistics

Guardian allows to display detailed statistics on the different clients that have queried the server. Here below is an example of the command result limited to 10 entries:

```
DNS Blast> show clients limit=10
```

Query	Client IP	View	Last-used	C-hit	C-miss	hit-Q-sz	miss-Q-sz	hit-A-sz
miss-A-sz	recurs-time		invalid-Q	restricted	block-Q	NOERROR	NXDOMAIN	
SERVFAIL	rpz-hit	trigger-hit						
235213	10.0.200.1	0	238	12891	222322	567KB	9571KB	1109KB
20229KB	16536k	0/0/0/0/0/0/0/0			0	8684	218595	
6502	0	98						
54364	10.10.13.37	0	2	51384	2980	1736KB	106KB	2763KB
278KB	181434	0/0/0/0/0/0/0/0			0	54249	111	
4	0	0						
38674	10.10.13.39	0	2	17684	20990	650KB	796KB	1512KB
2002KB	753664	0/0/0/0/0/0/0/0			0	37082	1548	
16	0	0						
28925	10.0.0.10	0	32	25077	3848	843KB	164KB	1923KB
361KB	44819	0/0/0/0/0/0/0/0			0	28280	645	
0	0	0						
27557	10.0.253.252	0	22	17584	9973	781KB	404KB	1893KB
1256KB	214175	0/0/0/0/0/0/0/0			0	27413	120	
24	0	10						
21213	10.0.0.17	0	15	17116	4097	648KB	167KB	1197KB
362KB	4199	0/0/0/0/0/0/0/0			0	11099	10113	
0	0	0						
17185	10.10.17.8	0	4	12531	4654	457KB	180KB	910KB
338KB	115822	0/0/0/0/0/0/0/0			0	16353	16	
0	0	0						
17064	10.10.13.41	0	63606	15624	1440	694KB	55577B	1446KB
143KB	66730	0/0/0/0/0/0/0/0			0	17058	2	
3	0	0						
16751	10.0.50.99	0	31	6799	9952	272KB	400KB	558KB

Managing Guardian Statistics

```

| 972KB | 91521 | 0/0/0/0/0/0/0/0 | 0 | 6149 | 10602 |
0 | 0 | 9 |
13596 | 10.0.30.254 | 0 | 1 | 3227 | 10369 | 129KB | 422KB | 253KB
| 1011KB | 2999 | 0/0/0/0/0/0/0/0 | 0 | 3758 | 9838 |
0 | 0 | 0
RCODE RETURNED:
RCODE | Total | Percent
NOERROR | 569609 | ( 63.87%)
SERVFAIL | 6673 | ( 0.75%)
NXDOMAIN | 308101 | ( 34.55%)

245 entries found (891816 total utilization)

search time: 0ms

```

The results of the command are described in the following table:

Table 57.6. Guardian client statistics

Data returned	Description
Query	Number of queries received from the client.
Client IP	IP address of the client.
View	Number of the view that answered the client query, from 0 to 7.
Last-used	Number of seconds since the client did query Guardian for the last time.
C-hit	Number of queries for which the answer was in the cache.
C-miss	Number of queries for which the answer was not in the cache.
hit-Q-sz	Total size (in bytes) of the queries which answer was in the cache.
miss-Q-sz	Total size (in bytes) of the queries which answer was not in the cache.
hit-A-sz	Total size (in bytes) of the queried answers that were in the cache.
miss-A-sz	Total size (in bytes) of the queried answers that were not in the cache.
recurs-time	Time needed to get the answer from the local recursive server, in milliseconds.
invalid-Q	<p>Number of invalid queries received from the client. This information follows the format <i>n/n/n/n/n/n/n/n</i> where <i>n</i> is the number. Each / separates a specific case:</p> <p>N/<i>n/n/n/n/n/n/n</i> is the Total sum of invalid queries received.</p> <p><i>n</i>/N/<i>n/n/n/n/n/n/n</i> is an Invalid class (IN).</p> <p><i>n/n</i>/N/<i>n/n/n/n/n/n/n</i> is an Overflow (wrong size).</p> <p><i>n/n/n</i>/N/<i>n/n/n/n/n/n/n</i> is a Discrepancy on the number of questions.</p> <p><i>n/n/n/n</i>/N/<i>n/n/n/n/n/n/n</i> is a Wrong opcode (use of reserved code).</p> <p><i>n/n/n/n/n</i>/N/<i>n/n/n/n/n</i> is an Invalid UDP source.</p> <p><i>n/n/n/n/n/n</i>/N/<i>n/n/n/n/n</i> is an IP SRC and IP DST are the same.</p> <p><i>n/n/n/n/n/n/n</i>/N/<i>n/n/n/n/n</i> is an IP SRC address is a broadcast address.</p> <p><i>n/n/n/n/n/n/n/n</i>/N is the number of packets of invalid size received.</p>
restricted	Action applied on the client: <i>any</i> , <i>none</i> , <i>quarantine</i> or <i>block</i> .
blocked-Q	Number of queries received and blocked by the server.
NOERROR	Number of queries returning a NOERROR error code.
NXDOMAIN	Number of queries returning a NXDOMAIN error code.
SERVFAIL	Number of queries returning a SERVFAIL error code.
rpz-hit	Number of times the client matched an RPZ entry.
trigger-hit	Number of hits triggered by the client.

Note that some of the client data is used to generate analytics available in the GUI. For more details, refer to the section [Displaying Guardian Analytics Tops](#).

To display Guardian client statistics

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. Make sure the client data retrieval is enabled. For more details, refer to the section [Enabling Guardian Client Statistics](#).
4. To display Guardian client statistics, type in the following command:

```
show clients
```

5. By default, the command displays 50 clients, in particular, the ones with the highest number of queries.

To limit the number of clients to be displayed, type in the following command:

```
show clients limit=<number>
```

6. To sort the entries by the type of data returned, described in the table [Guardian client statistics](#), type in the following command:

```
show clients order=<data-type>
```

7. To display specific client statistics, as described in the table [Guardian client statistics](#), use the following command with one or several of the filters from the table below:

```
show clients <filter>
```

Table 57.7. Available filters for Guardian client statistics

Filter	Description
used (> < = !=) <number-of-queries>	Filter the results using the number of queries received from the client.
view (> < = !=) <view-number>	Filter the results using the number of a DNS view .
lastused (> < = !=) <duration>	Filter the results using the last time the client queried the DNS server, in seconds.
C-miss (> < = !=) <number-of-queries>	Filter the results using the number of queries for which the answer was not in the cache.
C-hit (> < = !=) <number-of-queries>	Filter the results using the number of queries for which the answer was in the cache.
(<IP-address>/<prefix>)	Filter the results using the client IP address. For a range of addresses, use the format <ip-address>/<prefix>.
restricted (= !=) <restriction>	none: Display all the clients with no restriction at all. any: Display all the clients with any of the restrictions described below. ratelimit: Display all the clients that armed a trigger with the action <i>Ratelimit</i> . quarantine: Display all the clients that armed a trigger with the action <i>Quarantine</i> . quarantine_redirect: Display all the clients that armed a trigger with the action <i>Quarantine redirect</i> . block: Display all the clients that armed a trigger with the action <i>Block</i> . log: Display all the clients that armed a trigger with the action <i>Log only</i> .

You can use, at the same time, the order and limit parameters with one or several filters. For more details on trigger actions, refer to the section [Managing Triggers](#).

Resetting Guardian Client Statistics

You can reset client statistics at any time, thus bringing the counters of your choice to 0. This can prove useful for a real-time analysis of the traffic.

You can select the statistics to be reset using specific search parameters. For more details, refer to the table [Available filters for Guardian client statistics](#).

To reset Guardian client statistics

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To reset Guardian client statistics, use the following command:

```
reset clients
```

Clearing Guardian Client Statistics Manually

You can delete client statistics at any time. **Keep in mind that this action is irreversible.**

You can select the statistics to be cleared using specific search parameters. For more details, refer to the table [Available filters for Guardian client statistics](#).

To clear Guardian client statistics via CLI

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To clear Guardian client statistics, use the following command:

```
clear clients
```

To clear Guardian client statistics from the GUI

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Tick the Guardian server for which you want to flush clients.
3. In the menu, select **Edit > Command > Flush clients**. The wizard **Flush Guardian clients** opens.
4. In the field **IP address or network address (CIDR)**, specify the IP address of the client you want to flush or a network address to flush all the clients part of this network.

To flush all clients, type in `.*`

5. Click on **OK** to complete the operation. The wizard closes.

Clearing Guardian Client Statistics Automatically

You can use a set of parameters in Guardian configuration to automate the client statistics clearing process. **Keep in mind that this action is irreversible.**

Three parameters allow to define a maximum number of entries to keep in the client statistics, when to start the automatic clearing of the client statistics and the maximum number of remaining entries after the clearing. When the maximum percentage of entries is reached, the clearing commands are automatically executed in ascending order until the number of remaining queries equals or is lower than the minimum percentage of entries. You can configure up to 10 clearing commands matching the filters of your choice.

The automated clearing is configured using the parameters below:

- **Max client entries** sets the maximum of entries that you want to store in the cache. Its default value depends on the memory of each appliance:

Table 57.8. Max client entries default values

SOLIDserver model	Max client entries default value
SDS-550	20000
SDS-1100	75000
SDS-2200	500000
SDS-3300 & Blast series	1000000

- **Max clean client percent** sets the percentage of *Max client entries* that is to be reached to launch the execution of the clearing commands *Clean client<0-9>.cmd*.
- **Min clean client percent** sets the percentage of *Max client entries* to keep in the cache after executing the clearing commands *Clean client<0-9>.cmd*.

Once *Max clean client percent* is reached, the cleaning command *Clean client0.cmd* is executed. If the cache still contains more entries than the value of *Min clean client percent*, the command *Clean client1.cmd* is executed. All the commands *Clean client<0-9>.cmd* are executed in order until the percentage of entries remaining in the cache matches or is lower than the value of *Min clean client percent*.

- **Clean client<0-9>.cmd** parameters set 10 different clearing commands that respect the clearing filters of your choice.

The clearing commands are executed in order, but only if:

- The parameter *Clean client<0-9>.cmd* has a value.
- The number of entries left in the cache is higher than the value set with *Min clean client percent*.

By default, only the first two clearing commands *Clean client0.cmd* and *Clean client1.cmd* are configured.

To automate Guardian client statistics clearing

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **[EDIT]**. The wizard **Options configuration** opens.

4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To configure the clearing commands, in the fields **Clean client<0-9>.cmd**, use the filters detailed in the table [Available filters for Guardian client statistics](#).
7. In the field **Max client entries**, specify the maximum number of entries in Guardian cache.
8. In the field **Max clean client percent**, specify the percentage of the maximum number of client entries that will start the automatic clearing of client statistics.
9. In the field, **Min clean client percent**, specify the percentage of the maximum number of client entries that you want to keep in the client statistics after the clearing.
10. Click on **OK** to complete the operation. The wizard closes.

Monitoring Guardian from the GUI

In addition to CLI management, the GUI provides a selection of graphs on the server properties page as well as detailed analytics on clients queries.

Note that you must make sure that client statistics are enabled. For more details, refer to the section [Enabling Guardian Client Statistics](#).

Monitoring Guardian Statistics

From the GUI, you can monitor the Guardian cache statistics. When Guardian is configured and enabled, the properties page of the local EfficientIP DNS server provides extra panels containing charts dedicated to Guardian. Note that:

- All the information returned in these panels is based on some of the server statistics available in CLI, the data used is indicated after each description. For more details, refer to the table [Guardian server statistics](#).
- All the panels contain at least one chart. The y-axis of these charts indicates the unit, the axis scale and unit prefix depend on the period selected and maximum value displayed. Following the standard ISO 80000-1, all the y-axis units can have no prefix or any SI prefix such as: *m* (milli), *k* (kilo) or *M* (mega).

Keep in mind that you can zoom in and out of the charts or decide the period and data to display. For more details refer to the section [Charts](#).

To display Guardian statistics panels

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. In the column **DNS Guardian**, the server is marked **Enabled**.
3. At the end of the line of the server of your choice, click on **ⓘ**. The server properties pages opens.

Keep in mind that these panels can be managed like a gadget and displayed on any dashboard. For more details, refer to the section [Assigning a Gadget from a Resource Properties Page](#).

More detailed statistics and options are available via CLI. For more details, refer to the sections [Managing Guardian server Statistics](#) and [Managing Guardian Client Statistics](#).

Each of these panels is named **Guardian - <data>** and described below:

Guardian - Cache Size (Bytes)

A chart displaying Guardian cache size evolution over time, in bytes. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Cache size*.

Guardian - Cache Hits/Misses

A chart displaying the ratio, in percent, of Guardian answered and unanswered queries. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Cache hit rate* and *Cache miss rate*.

Guardian - Hits in Quarantine

A chart displaying the number of queries per seconds that Guardian answered to clients impacted by the trigger action *Quarantine* or *Quarantine redirect*. For more details, refer to the section [Managing Triggers](#).

Based on: *Cache hit in quarantine mode*.

Guardian - Cache Statistics

A chart displaying Guardian cache statistics, in queries per second, for all DNS traffic. The results include traffic in quarantine and rescue mode. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Cache hit*, *Cache missed*, *Cache missed but exists (expired)* and *Cache missed, never queried*.

Guardian - Quarantine statistics

A chart displaying Guardian cache statistics, in queries per second, for clients impacted by the trigger action *Quarantine* or *Quarantine redirect*. For more details, refer to the section [Managing Triggers](#).

Based on: *Cache hit in quarantine mode*, *Cache missed in quarantine mode*, *Cache missed but exists (expired) in quarantine mode* and *Cache missed, never queried in quarantine mode*.

Guardian - Rescue mode statistics

A chart displaying Guardian cache hit statistics, in queries per second, for clients in rescue mode. For more details regarding the rescue mode, refer to the section [Managing DNS Guardian Rescue Mode](#).

Based on: *Cache hit in rescue mode*, *Cache missed in rescue mode*, *Cache missed but exists (expired) in rescue mode* and *Cache missed, never queried in rescue mode*.

Guardian - Blocked traffic (Queries)

A chart displaying the cumulated number of queries, per second, generated by all the clients impacted by the trigger action *Block*. For more details, refer to the section [Managing Triggers](#).

Based on: *Blocked queries*.

Guardian - Tracked clients

A chart displaying the number of client IP addresses impacted by the trigger action *Quarantine* or *Quarantine redirect*. For more details, refer to the section [Managing Triggers](#).

Based on: *Client size*.

Guardian - Sent/Received traffic (Queries)

A chart displaying the number of queries, per second, sent and received by Guardian. For more details, refer to the section [Sending the Cache Content to Another Guardian server](#).

Based on: *Sent DNS answers* and *Received DNS queries*.

Guardian - Sent/Received traffic (Bytes)

A chart displaying the overall traffic, in bytes per second, generated by the queries sent and received by Guardian. For more details, refer to the section [Sending the Cache Content to Another Guardian server](#).

Based on: *Sent DNS answers (bytes)* and *Received DNS queries (bytes)*.

Guardian - Answers in Quarantine/Rescue Mode (Queries)

A chart displaying the number of queries, per second, answered by Guardian for clients impacted by the trigger action *Quarantine* or *Quarantine redirect* or clients in rescue mode. For more details, refer to the sections [Managing Triggers](#) and [Managing DNS Guardian Rescue Mode](#).

Based on: *Sent DNS answers in quarantine mode* and *Sent DNS answers in rescue mode*.

Guardian - Answers in Quarantine/Rescue Mode (Bytes)

A chart displaying the overall traffic generated by the queries answered by Guardian for clients impacted by the trigger action *Quarantine* or *Quarantine redirect* and clients in rescue mode, in bytes per second. For more details, refer to the sections [Managing Triggers](#) and [Managing DNS Guardian Rescue Mode](#).

Based on: *Sent DNS answers in quarantine mode (bytes)* and *Sent DNS answers in rescue mode (bytes)*.

Guardian - Invalid traffic (Queries)

A chart displaying the total number of invalid queries that were dropped by Guardian. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Received Invalid DNS queries*.

Guardian - Invalid traffic (Bytes)

A chart displaying the traffic generated by all the invalid queries that were dropped by Guardian, in bytes. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Received Invalid DNS queries (bytes)*.

Guardian - Cache sharing statistics

A chart displaying the number of entries sent and received via shared Guardian cache, in queries per second. For more details, refer to the section [Sharing the Cache of Several Guardian servers](#).

Based on: *DNS entries received via shared cache* and *DNS entries sent via shared cache*.

Guardian - Cache sharing statistics (Bytes)

A chart displaying the traffic generated by all the entries sent and received via shared Guardian cache, in bytes. For more details, refer to the section [Sharing the Cache of Several Guardian servers](#).

Based on: *DNS entries received via shared cache (bytes)* and *DNS entries sent via shared cache (bytes)*

Guardian - Return codes statistics

A chart displaying the number of return codes sent by Guardian, in queries per second. For more details, refer to the chapter [Managing Guardian Cache](#).

Based on: *Sent NOERROR rcode, Sent FORMERR rcode, Sent SERVFAIL rcode, Sent NXDOMAIN rcode, Sent NOTIMP rcode, Sent REFUSED rcode, Sent YXDOMAIN rcode, Sent YXRRSET rcode, Sent NXRRSET rcode, Sent NOTAUTH rcode and Sent NOTZONE rcode.*

Guardian - Triggers

A chart displaying the number of times each of the first 8 available triggers were armed. The chart is generated using the position of each trigger, therefore even if you edit the trigger name, action, period and/or threshold, the graph keeps evolving but might not always have the same meaning, especially if you duplicate a policy as it resets the trigger positions. For more details, refer to the sections [Managing Guardian Policies](#) and [Managing Triggers](#).

Based on: *Number of time trigger #0 was armed, Number of time trigger #1 was armed, Number of time trigger #2 was armed, Number of time trigger #3 was armed, Number of time trigger #4 was armed, Number of time trigger #5 was armed, Number of time trigger #6 was armed and Number of time trigger #7 was armed.*

Guardian - Recursive queries latency

A chart displaying the recursive query latency distribution of Guardian, in queries per second. The latency is divided into 6 periods and ranges from less than 10 milliseconds to more than 1600 milliseconds. Clients traffic rate can be limited to 10 queries per ms using the trigger action *Ratelimit*. For more details, refer to the section [Managing Triggers](#).

Based on: *Recursive queries with latency < 10ms, Recursive queries with latency >= 10ms and < 100ms, Recursive queries with latency >= 100ms and < 500ms, Recursive queries with latency >= 500ms and < 800ms, Recursive queries with latency >= 800ms and < 1600ms and Recursive queries with latency >= 1600ms.*

Guardian - Rate limited traffic (Queries)

A chart displaying the number of queries dropped by Guardian because source clients reached the configured query limit rate (of 100 QPS), in queries per second. Clients traffic rate can be limited to 10 queries per ms using the trigger action *Ratelimit*. For more details, refer to the section [Managing Triggers](#).

Based on: *Ratelimited queries.*

Monitoring Guardian Analytics From the GUI

From the GUI, you can monitor Guardian clients data via detailed analytics, for one or several servers, as long as the related analytics collection is enabled.

Note that when Guardian is enabled and configured:

- Standard *DNSTOP* Analytics are no longer available. For more details, refer to the section [Displaying the DNS Analytics](#).
- The sampling *Periodicity (min.)*, in the panel **Analytics** and the related wizard, is no longer visible as it is not configurable.

Making Sure Guardian Analytics are Collected

To display Guardian analytics, you must make sure that DNS analytics collection is enabled.

To make sure Guardian Analytics are collected

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. Filter the column **Guardian/GSLB** to only display *Enabled* servers.

3. In the column **Name**, right-click over the server of your choice. The contextual menu opens.
4. Click on **Properties**. The server properties page opens.
5. In the panel **DNS Analytics**, click on **EDIT**. The wizard **Configure DNS Analytics** opens.
6. Tick the box **Enable analytics collection**.
7. Click on **OK** to complete the operation. The report opens and closes. The server properties pages is visible again. In the panel **Analytics**, the option *Enable analytics collection* is set to yes.

Displaying Guardian Analytics Tops

When Guardian is configured and enabled, the local DNS EfficientIP server provides dedicated tops analyzing Guardian data on the page *Analytics*.

Guardian tops allows to display, for a **time window** that can range from 5 minutes to 1 week, every **5-minute samples** of the traffic in which **a metric has been measured**.

Each top provides:

- A set of dedicated filters that allow to display a precise view of the traffic, all detailed in the image below.
- Specific columns.
- Three different views of the data returned.

For instance, as in the figure below, you can use the *Overview* to display the total number of queries received by Guardian during the 3 hours that precede a certain time and date:

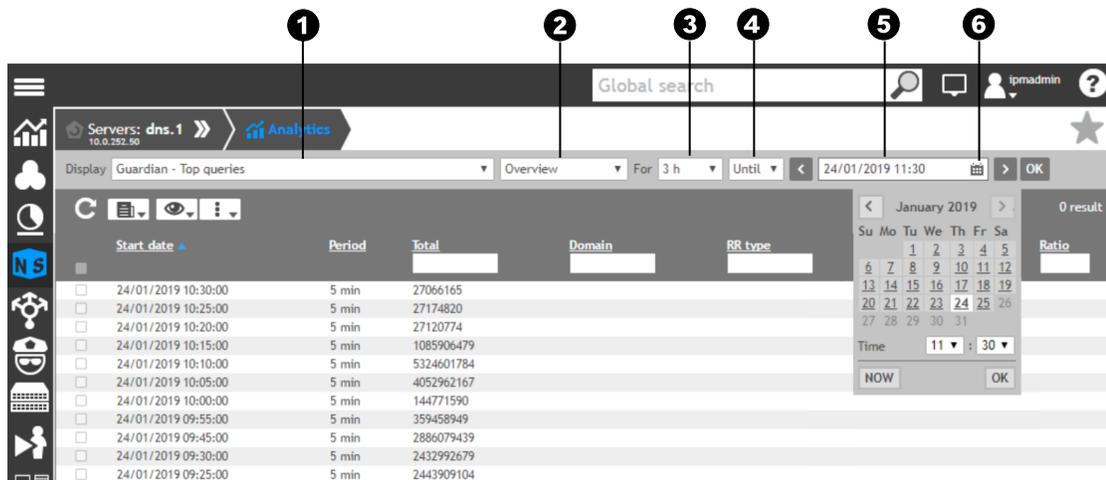


Figure 57.1. Guardian Analytics filters

- 1 The drop-down list **Display** allows to select the analytics of your choice. If you are displaying the page *Analytics* of a Guardian server, only *Guardian - Top <data>* are listed.
- 2 There are three displays available: **Overview**, **Detailed view** and **Consolidated view**. No matter the selected display, the number of entries listed depend on the time window you set using the next three fields. For more details, refer to the section [Using the Views to Analyse a Specific Guardian Top](#).
- 3 The drop-down list **For** allows to choose a time window for the analytics, ranging from 5 minutes to a week. The available durations depend on the view you choose.

- 4 This drop-down list allows to set the bounds of your search: either **Since**, to display the Top data for the selected duration from the date and time; or **Until**, to display the Top data for the selected duration until the date and time.
- 5 This field allows to specify the date and date of your choice. The buttons  and  on each side of the field allow to edit the field content based on the duration selected in the list *For*.
- 6 The icon  allows to open the time and date constructor and browse the calendar to select a date and set a specific time using the drops-down lists. Once you selected both date and time, click on OK to refresh the field and display Guardian analytics.

From the page *Analytics*, you can display all the available tops.

To display Guardian Analytics data

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. To launch a cross-servers analysis, go to the step 4.
3. To launch an analysis on a specific server, at the end of the line of the server of your choice, click on . The server properties pages opens.
4. In the breadcrumb, on the right of **All servers**, click on  to display additional pages.
5. Click on **Analytics**. The page opens.
6. In the drop-down list **Display**, only Guardian data is available and named as follows: *Guardian - Top <data>*.

All the columns can be used to filter and sort the data, except for the columns *Start date* and *Period*, where filtering is not possible. Depending on the top and view selected, some columns may be empty or not displayed.

Table 57.9. Guardian analytics columns

Column	Description
Server	The name of the Guardian server(s) if you launched a cross-servers analysis. This column may not appear if the server is already selected and displayed in the breadcrumb.
Start date	In the <i>Overview</i> and <i>Detailed view</i> : the start date and time of the 5-minute sample <i>Period</i> .
	In the <i>Consolidated view</i> : the start date and time of the first sample <i>Period</i> , within the analytics time window, during which the metric has been measured.
Period	The sample period, either detailed (5 minutes) or consolidated (multiple of 5 minutes), during which metrics have been measured.
Client's IP address	The IP address of the client querying Guardian. Depending on the top and view selected, this column may be empty or not displayed.
Domain	The name of the domain queried. Depending on the top and view selected, this column may be empty or not displayed.
RR type	The record returned when the domain was queried: <i>A</i> , <i>AAAA</i> , <i>PTR</i> ... Depending on the top and view selected, this column may be empty or not displayed.
<metric>	<i>Queries</i> : the number of queries received by the server on the sample period, either from a <i>Client's IP address</i> or for a specific <i>Domain</i> and <i>RR Type</i> .
	<i>Time (ms)</i> : the recursion time, in milliseconds, of the queries received by the server on the sample period, either from a <i>Client's IP address</i> or for a specific <i>Domain</i> and <i>RR Type</i> .
	<i>Size (bytes)</i> : the size, in bytes, of the queries received by the server on the sample period, either from a <i>Client's IP address</i> or for a specific <i>Domain</i> and <i>RR Type</i> .

Column	Description
	<i>Trigger hits</i> : the number of times, in the sample period, a <i>Client's IP address</i> armed a trigger.
Total	The value of all the metrics measured during the same sample period, regardless of the <i>Client's IP address, Domain or RR Type</i> .
Ratio	The proportion of the metric value compared to the total value within the same period, in percent.

Available Guardian Analytics Tops

Guardian allows you to display detailed measurements of the server traffic, in number of queries, size or recursion time. These metrics either return information based on the IP address of a client or domain/record queried.

Each of these analytics is named **Guardian - Top <data>** and described below.

All the information returned in these tops is based on some of the server statistics available in CLI, the data used is indicated after each description. For more details, refer to the tables [Guardian client statistics](#) and [Guardian analytics columns](#).

Guardian - Top clients generating dropped queries (BLOCK)

The clients, whose query was received and dropped because a trigger with the action *Block* was armed. For more details, refer to the section [Managing Triggers](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *blocked-Q*.

Guardian - Top clients querying records missing from the cache

The clients querying records not in the cache. For more details, refer to the sections [Managing Guardian Client Statistics](#) and [Managing Guardian server Statistics](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *C-miss*.

Guardian - Top clients querying records present in the cache

The clients querying records in the cache. For more details, refer to the sections [Managing Guardian Client Statistics](#) and [Managing Guardian server Statistics](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *C-hit*.

Guardian - Top clients generating recursive incoming traffic

The total size of the answers queried by clients that were not in the cache and were received from the local recursive server. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Size (bytes)* by *Client's IP address*.

Based on: *miss-A-sz*.

Guardian - Top clients generating recursive outgoing traffic

The total size of the queries whose answer was not in the cache and that were sent to the local recursive server. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Size (bytes)* by *Client's IP address*.

Based on: *miss-Q-sz*.

Guardian - Top clients

The clients querying Guardian. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *Query*.

Guardian - Top clients receiving NOERROR rcodes

The clients whose queries were answered with the rcode NOERROR. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *NOERROR*.

Guardian - Top clients receiving NXDOMAIN rcodes

The clients whose queries were answered with the rcode NXDOMAIN. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *NXDOMAIN*.

Guardian - Top recursion time by client

The clients for whom the cumulative queries took the longest time to be answered by the local recursive server, in milliseconds. For more details, refer to the section [Managing Guardian Client Statistics](#).

Metrics: *Time (ms)* by *Client's IP address*.

Based on: *recurs-time*.

Guardian - Top clients receiving SERVFAIL rcodes

The clients whose queries were answered with the rcode SERVFAIL. For more details, refer to the sections [Managing Guardian Client Statistics](#) and [Ignoring the SERVFAIL Error Message Differences](#).

Metrics: *Queries* by *Client's IP address*.

Based on: *SERVFAIL*.

Guardian - Top clients triggering policies

The number of times a client query armed a trigger and, therefore, triggered an action. For more details, refer to the sections [Managing Guardian Client Statistics](#) and [Managing Triggers](#).

Metrics: *Trigger Hits* by *Client's IP address*.

Guardian - Top queries of records missing from the cache

The queries sent to Guardian for records not in the cache. For more details, refer to the section [Managing Guardian server Statistics](#).

Metrics: *Queries*.

Based on: *C-miss*.

Guardian - Top queries

The queries sent to Guardian. For more details, refer to the section [Managing Guardian server Statistics](#).

Metrics: *Queries*.

Based on: *Used*.

Guardian - Top SLD creating new cache entries

The queries for a sub-level domain name that has not been queried before. For more details, refer to the section [Managing Guardian server Statistics](#).

Metrics: *Queries by Domain* and *RR Type*.

Based on: *Query* and *first-used*.

Guardian - Top queries returning NXDOMAIN rcodes

The queries returned with the rcode NXDOMAIN. For more details, refer to the section [Managing Guardian server Statistics](#).

Metrics: *Queries*.

Based on: *RCODE*.

Guardian - Top recursion time by FQDN

The recursion of all queries, one fully qualified domain name at a time. For more details, refer to the section [Managing Guardian server Statistics](#).

Metrics: *Time (ms) by Domain* and *RR Type*.

Based on: *R-time*.

Guardian - Top queries returning SERVFAIL rcodes

The queries returned with the rcode SERVFAIL. For more details, refer to the sections [Managing Guardian server Statistics](#) and [Ignoring the SERVFAIL Error Message Differences](#).

Metrics: *Queries*.

Based on: *RCODE*.

They all provide three different ways to display the analytics within the time window set in the drop-down list *For*.

Note that the analytics time window can range from 5 minutes to 6 hours for *Detailed view* and *Consolidated view* or 1 hour to 1 week for the *Overview*. For more details, refer to the section [Using the Views to Analyse a Specific Guardian Top](#).

Using the Views to Analyze a Specific Guardian Top

As described in the section [Displaying Guardian Analytics Tops](#), each Guardian Top offers three different views of the data for the same time window:

1. **Overview** returns a simplified display of the data for each 5-minute sample matching the time window set in your filters. With the column *Total*, you can identify peaks in the traffic and get more information with the consolidated or detailed view.

Server	Start date	Period	Domain	RR type	Queries	Total	Ratio
netdns.intra	24/01/2019 10:50:00	5 min				1013	
netdns.intra	24/01/2019 10:55:00	5 min				511	
netdns.intra	24/01/2019 11:00:00	5 min				331	
netdns.intra	24/01/2019 11:05:00	5 min				474	
netdns.intra	24/01/2019 11:10:00	5 min				417	
netdns.intra	24/01/2019 11:15:00	5 min				604	

Figure 57.2. Overview of Guardian Analytics

2. **Consolidated view** focuses on the client's IP address or on the domain and RR type, depending on the selected Top. When a client IP address or queried domain/record is repeatedly identified over the different 5-minute samples matching the time window set in your filters, these sampling periods add up in the column *Period*.

Server	Start date	Period	Domain	RR type	Queries	Total	Ratio
netdns.intra	24/01/2019 10:50:00	120 min	loginsight.intranet		34	11702	0.29%
netdns.intra	24/01/2019 10:50:00	120 min	freenas.intranet		24	11702	0.21%
netdns.intra	24/01/2019 10:50:00	120 min	vcenter.intranet		213	11702	1.82%
netdns.intra	24/01/2019 10:50:00	120 min	loginsight		34	11702	0.29%
netdns.intra	24/01/2019 10:50:00	70 min	play.google.com		123	6763	1.82%
netdns.intra	24/01/2019 10:50:00	70 min	www.google.com		110	7265	1.51%

Figure 57.3. Consolidated view of Guardian Analytics

3. **Detailed view** focuses on a sampling period. For each 5-minute sampling period, this view lists the number of times one IP address or a domain and RR type matching your filters have been identified.

Server	Start date	Period	Domain	RR type	Queries	Total	Ratio
netdns.intra	24/01/2019 10:50:00	5 min	www.clipart-fr.com	A	5	1013	0.49%
netdns.intra	24/01/2019 10:50:00	5 min	notifications.google.com	A	4	1013	0.39%
netdns.intra	24/01/2019 10:50:00	5 min	clients4.google.com	A	8	1013	0.79%
netdns.intra	24/01/2019 10:50:00	5 min	adservice.google.fr	A	5	1013	0.49%
netdns.intra	24/01/2019 10:50:00	5 min	www.gstatic.com	A	4	1013	0.39%
netdns.intra	24/01/2019 10:50:00	5 min	cm.g.doubleclick.net	A	4	1013	0.39%

Figure 57.4. Consolidated view of Guardian Analytics

You can select the **Overview** on the *Guardian - Top queries returning NXDOMAIN rcodes* to display a 6-hour window around a specific time and date. Sorting the data with the column *Total* allows to pinpoint the 5-minute samples during which Guardian handled an unusual peak of potentially suspicious queries.

To display a Guardian Top Overview

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on **ⓘ**. The server properties pages opens.
3. In the breadcrumb on the right of the server name, click on **»** to display additional pages.

4. Click on **Analytics**. The page refreshes.
5. In the drop-down list **Display**, select *Guardian - Top <data>* of your choice.
6. In the next drop-down list, select *Overview*.
7. In the drop-down list **For**, select *1 h, 3 h, 6 h, 1 day* or *1 week*.
8. In the next drop-down list, select **Until** to use the time and date and duration you set as an end point for the search results or **Since** to set them as a starting point for the search.

Note that if you select *Until* and specify a time such as *xx:00* or *xx:05*, this particular time is excluded from the result because it is the starting point of a 5-minute sampling going from *xx:00* to *xx:05* or *xx:05* to *xx:10*. To include *xx:05* to the result, specify *xx:06*.

9. In the last field, specify the date and time that suit your needs. You can either type in the date and time following the format *dd/mm/yyyy hh:mm*, or use the time and date constructor:
 - a. To open the time and date constructor, click on  or double-click in the field.
 - b. To set the date, click on  and  to browse the months and select the date of your choice. The date selected is highlighted in white.
 - c. To set the **Time**, select the hour in the first drop-down list and the minutes in the second.
 - d. When you selected both date and time, click on . The constructor closes and the date and time you set are displayed in the field.

Note that the arrows  and  on both sides of the date and time field allow to edit the field content based on the duration set in the list **For**. Click on either to add or subtract 5 minutes, an hour, six hours or even a day to the date and time displayed.

10. At the end of the line, click on . The entries matching the selected Top and filters are displayed.

Only the columns *Start date*, *Period* and *Total* contain data.

You can also select the **Consolidated view** on the *Guardian - Top clients triggering policies* to display a 6-hour window around a specific time and date. Sorting the data with the column *Period* allows to highlight the clients that repeatedly triggered actions over time.

To display a Guardian Top Consolidated view

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The server properties pages opens.
3. In the breadcrumb on the right of the server name, click on  to display additional pages.
4. Click on **Analytics**. The page refreshes.
5. In the drop-down list **Display**, select *Guardian - Top <data>* of your choice.
6. In the next drop-down list, select *Consolidated view*.
7. In the drop-down list **For**, select *5 min, 30 min, 1 h, 3 h* or *6 h*.
8. In the next drop-down list, select **Until** to use the time and date and duration you set as an end point for the search results or **Since** to set them as a starting point for the search.

Note that if you select *Until* and specify a time such as *xx:00* or *xx:05*, this particular time is excluded from the result because it is the starting point of a 5-minute sampling going from *xx:00* to *xx:05* or *xx:05* to *xx:10*. To include *xx:05* to the result, specify *xx:06*.

9. In the last field, specify the date and time that suit your needs. You can either type in the date and time following the format *dd/mm/yyyy hh:mm*, or use the time and date constructor:
 - a. To open the time and date constructor, click on  or double-click in the field.
 - b. To set the date, click on  and  to browse the months and select the date of your choice. The date selected is highlighted in white.
 - c. To set the **Time**, select the hour in the first drop-down list and the minutes in the second.
 - d. When you selected both date and time, click on . The constructor closes and the date and time you set are displayed in the field.

Note that the arrows  and  on both sides of the date and time field allow to edit the field content based on the duration set in the list **For**. Click on either to add or subtract 5 minutes, an hour, six hours or even a day to the date and time displayed.

10. At the end of the line, click on . The entries matching the selected Top and filters. All the columns contain data.

Finally, you can select the **Detailed view** on the *Guardian - Top clients* to display a 6-hour window around a specific time and date. Sorting the data with the column *Client's IP address* allows to display every 5-minute period during which a certain client queried Guardian.

To display a Guardian Top Detailed view

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the server or smart architecture of your choice, click on . The server properties pages opens.
3. In the breadcrumb on the right of the server name, click on  to display additional pages.
4. Click on **Analytics**. The page refreshes.
5. In the drop-down list **Display**, select *Guardian - Top <data>* of your choice.
6. In the next drop-down list, select *Detailed view*.
7. In the drop-down list **For**, select *5 min, 30 min, 1 h, 3 h* or *6 h*.
8. In the next drop-down list, select **Until** to use the time and date and duration you set as an end point for the search results or **Since** to set them as a starting point for the search.

Note that if you select *Until* and specify a time such as *xx:00* or *xx:05*, this particular time is excluded from the result because it is the starting point of a 5-minute sampling going from *xx:00* to *xx:05* or *xx:05* to *xx:10*. To include *xx:05* to the result, specify *xx:06*.

9. In the last field, specify the date and time that suit your needs. You can either type in the date and time following the format *dd/mm/yyyy hh:mm*, or use the time and date constructor:
 - a. To open the time and date constructor, click on  or double-click in the field.
 - b. To set the date, click on  and  to browse the months and select the date of your choice. The date selected is highlighted in white.
 - c. To set the **Time**, select the hour in the first drop-down list and the minutes in the second.

- d. When you selected both date and time, click on . The constructor closes and the date and time you set are displayed in the field.

Note that the arrows  and  on both sides of the date and time field allow to edit the field content based on the duration set in the list **For**. Click on either to add or subtract 5 minutes, an hour, six hours or even a day to the date and time displayed.

10. At the end of the line, click on . The entries matching the selected Top and filters are displayed. All the columns contain data.

Chapter 58. Managing Guardian Protection

Malicious operations on the DNS infrastructure, such as DDoS, DNS tunneling or data exfiltration attacks, usually result in overloaded servers and service loss for legitimate clients. You can use the information provided by the Guardian client statistics to identify the different types of threats.

For instance, a high *C-miss* value might indicate that a client keeps querying names that are not cached, thus increasing the load on the recursive server. You can also correlate an important *miss-A-sz* to an encapsulated data transfer, taking advantage of the number of characters allowed in the packets. As for the *miss-Q-sz* counter, a large size of sent queries might be due to data exfiltration on the same principle. Other counters like *recurs_time* or *invalid-Q* are also good indicators in the case of a sloth or invalid queries attack.

With Guardian:

- If your recursive server is **under DDoS attack**, a **Rescue mode temporarily provides best-effort delivery** and prevents your server from being overloaded or victim of cache-poisoning. For more details, refer to the section [Managing DNS Guardian Rescue Mode](#).
- You can **configure up to 8 Guardian lists of domains**, either added manually or automatically updated, to gather specific metrics to combine with views and triggers. For more details, refer to the section [Managing Guardian Lists](#).
- You can **configure up to 8 Guardian views** that you can either associate to a transparent DNS proxy or to Guardian lists, in order to apply domain specific access policies. For more details, refer to the section [Managing Guardian Views](#).
- From the GUI, you can **create and configure Guardian policies** that are containers for triggers. For more details, refer to the section [Managing Guardian Policies](#)
- You can **target suspicious clients** and either monitor them, put them in quarantine, redirect them, limit their query rate or block them completely. For more details refer to the section [Managing Triggers](#).

Note that, **except for the Rescue mode, DNS Guardian protection requires that you enable and understand Guardian client statistics**. For more details, refer to the table [Managing Guardian Client Statistics](#).

Enabling Guardian Protection

By default, Guardian Protection is enabled when the service is started. For more details, refer to the section [service DNS Guardian](#). However it can be disabled to stop using Guardian cache to answer client queries directly but still update its content. For more details, refer to the section [Disabling Guardian Protection](#).

To enable Guardian protection

1. In the sidebar, go to  **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on . The properties page opens.
3. In the panel **Options**, click on [EDIT](#). The wizard **Options configuration** opens.

4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To enable DNS Guardian protection, in the field **Blast**, select *yes (1)*.
7. Click on **OK** to complete the operation. The wizard closes.

Managing Guardian Rescue Mode

Guardian provides SOLIDserver with a Rescue mode that allows to answer as many queries as possible if the local DNS server is unable to do so due to high traffic, as during DDOS attacks.

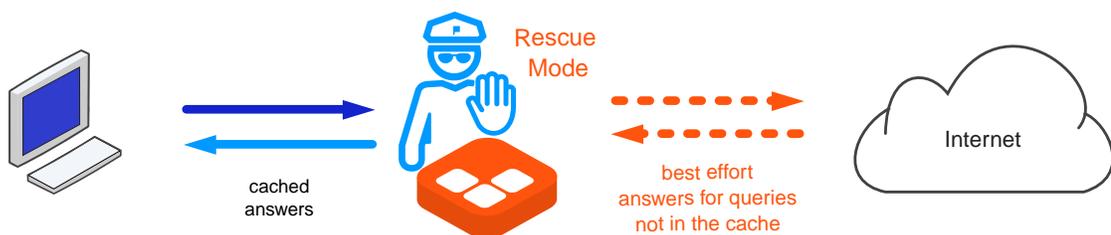


Figure 58.1. DNS Guardian's behavior in Rescue mode

In Rescue mode, Guardian buffers all the queries and behaves as follows:

1. If the queried record is in the cache, the response is immediate.
2. Usually, if the queried record is not cached yet, Guardian relays it to your local DNS server. In Rescue mode, the local DNS server offers a best-effort service to deliver answers to the clients. If the query is answered, it is cached by Guardian.
3. If the queried record has expired, Guardian sends it to the client with the TTL that you have set beforehand (300 seconds by default) to preserve the local DNS server and potentially avoid querying it altogether. Keep in mind that you can manually expire all or part of your cache entries.

You first need to [enable Rescue detection](#) so that the Rescue mode can be automatically triggered when the parameters set when you [configure Guardian Rescue mode](#) are met. You can also [force Guardian Rescue mode](#) and then stop it manually. If you do not want the Rescue mode to be triggered automatically, you can [disable Guardian Rescue detection](#).

Note that:

- For the Rescue mode to be available, Guardian Protection must be enabled. For more details, refer to the section [Enabling Guardian Protection](#).
- Guardian provides you with charts to visualize different Rescue Mode statistics on the properties page of the server. For more details, refer to the section [Monitoring Guardian Statistics from the GUI](#).

Enabling Rescue Detection

Rescue detection is enabled by default but can be disabled. For more details, refer to the section [Configuring Guardian Rescue Mode](#).

To enable Rescue detection

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To enable the Rescue mode detection, in the field **Rescue detection**, select *auto (1)*.
7. Click on **OK** to complete the operation. The wizard closes.

Configuring Guardian Rescue Mode

The rescue mode can be triggered for one of the three following reasons:

1. The server reaches a maximum number of queries per second sent from Guardian to the local DNS server: above the set value, Guardian switches to Rescue mode. Refer to *Rescue high rec packet* in the table below.
2. The server reaches both:
 - A minimum number of queries per second sent from Guardian to the local DNS server. Refer to *Rescue min rec packet* in the table below.
 - A percentage of these queries that are unanswered. This value acts as a threshold. Refer to *Rescue unanswered rate* in the table below.
3. The server receives a maximum number or percentage of SERVFAIL error messages per second from the local resolver. Refer to *Rescue servfail qps* or *Rescue ratio servfail* in the table below.

Table 58.1. Rescue mode parameters

Parameter	Default value	Description
Rescue time	5	Sets a check period (<i><value>*10s</i>) during which the inbound traffic, out-bound traffic and packet traffic between Guardian and the local DNS server are monitored to determine if a switch to Rescue mode is needed. Over this period, a check is performed every 10 seconds and the <i><value></i> sets the number of checks performed. By default, it is set to 5, meaning that over the last 50 seconds (5*10s), DNS Guardian checked 5 times every 10 seconds if the traffic met the triggering conditions that determine if a switch to Rescue mode is necessary.
Rescue ttl	300	Defines the TTL values of the cached records queried by clients that already expired once Guardian is in Rescue mode. It overwrites the actual record TTL, so we suggest that you set a low value to prevent clients from keeping an outdated value once Guardian turns off the Rescue mode.

Parameter	Default value	Description
Rescue min rec packet	5000	Sets the minimum number of queries per second received by the local DNS server that call for an extra check. If the number of queries set is met, Guardian checks the condition of the parameter <i>Rescue unanswered rate</i> to decide if it is necessary to switch to Rescue mode.
Rescue unanswered rate	10	Sets a percentage of unanswered queries by the local DNS server, based on the number of queries of <i>Rescue min rec packet</i> . If the percentage of unanswered queries is exceeded, Guardian switches to Rescue mode.
Rescue high rec packet	200000	Sets the maximum number of queries per second sent from Guardian to the local DNS server, above the value set, Guardian switches to Rescue mode.
Rescue servfail qps	50000	Sets the maximum number of queries per second answered by the local resolver with SERVFAIL error messages.
Rescue ratio servfail	90	Sets the maximum percentage of queries per second answered by the local resolver with SERVFAIL error messages.

You can set Rescue mode parameters to match your network needs.

To configure Rescue mode parameters

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To set a rescue mode monitoring period, in the field **Rescue time**, specify a check period: *<number-of-checks> * 10 seconds*.
7. To define the expired records TTL once in Rescue mode, in the field **Rescue ttl**, specify a number of seconds.
8. To trigger the Rescue mode when the number of queries received per second matches a specific number and the percentage of unanswered queries reaches a specific value,
 - in the field **Rescue min rec packet**, specify the number of queries received per second,
 - in the field **Rescue unanswered rate**, specify the percentage of unanswered queries.
9. To trigger the Rescue mode when the number of received queries per second reaches a specific number, in the field **Rescue high rec packet**, specify a number of queries.
10. To trigger the Rescue mode based on a maximum number of queries per second answered with SERVFAIL error messages, in the field **Rescue servfail qps**, specify a number of queries.
11. To trigger the Rescue mode based on a maximum percentage of queries per second answered with SERVFAIL error messages, in the field **Rescue ratio servfail**, specify a percentage.
12. Click on **OK** to complete the operation. The wizard closes.

Forcing Guardian Rescue Mode

When Rescue detection is enabled, you can manually force Guardian to switch to Rescue Mode, regardless of the threshold(s) you set in the parameters. For more details, refer to the section [Enabling Rescue Detection](#).

To switch to Rescue mode manually

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To manually switch Guardian to Rescue mode, in the field **Rescue detection**, select *forced (2)*.
7. To switch back to normal Rescue mode detection, which uses the threshold(s) you set in the parameters, in the field **Rescue detection**, select *auto (1)*.
8. Click on **OK** to complete the operation. The wizard closes.

Disabling Guardian Rescue Detection

Rescue detection is enabled by default but can be manually disabled, in which case, Guardian does not switch to Rescue mode if the traffic reaches the threshold(s) you set in the parameters.

To disable Guardian Rescue mode

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display options(s)**, select *using default values*.
6. To disable the Rescue mode detection, in the field **Rescue detection**, select *disabled (0)*.
7. Click on **OK** to complete the operation. The wizard closes.

Managing Guardian Lists

You can configure up to 8 lists of domains which can be used as whitelists, blacklists or any kind of restriction lists.

Guardian lists can be used to:

- Filter clients traffic querying specific domains depending on the view they are associated with. For more details, refer to the section [Filtering Guardian Views Using Lists](#).

- Set a Guardian trigger threshold using metrics of certain group of domains. For more details, refer to the section [Adding Triggers Relying on Lists Metrics](#).

From the GUI, you can set the [Guardian lists configuration parameters](#) to rename a Guardian list, update its content from an external source or save or clear a list when the service DNS Guardian restarts.

Via CLI, you can [display Guardian lists entries](#), [edit Guardian list content](#), [reset a Guardian list counter](#) and [clear guardian lists](#).

Displaying Guardian Lists

By default, all 8 Guardian lists are empty. You can add as many entries in the list of your choice either manually or automatically from a remote DNS zone. For more details, refer to the section [Editing a Guardian List Content](#).

You can display the entries present in all Guardian lists via CLI, as in the example below:

```
DNS Blast> show list

Used | * | = | EOL | List ID members | Name
-----|-----|-----|-----|-----|-----
0 | | = | 0 | 4 | *.0-2u.com
0 | | = | 0 | 4 | *.0-800horoscopes.com
0 | | = | 0 | 4 | *.0-apr-card-credit-4u.info
0 | | = | 0 | 4 | *.0-apr-credit-card-4u.info
0 | | = | 0 | 4 | *.0-forona-ptclean-0.com
0 | | = | 0 | 4 | *.0-holds-barred.com
0 | | = | 0 | 4 | *.0-p-0.xz.cn
0 | | = | 0 | 4 | *.0-pm.com
0 | | = | 0 | 0 | *.anothertrusteddomain.com
0 | | = | 0 | 2 | *.artemislist.gti.myantivirus.com
0 | | = | 0 | 2 | *.avqs.myantivirus.com
0 | | = | 0 | 2 | *.avqs.myantivirus.com
0 | | = | 0 | 2 | *.cloud.gti.myantivirus.com
0 | | = | 0 | 2 | *.cwl2.gti.myantivirus.com
0 | | = | 0 | 2 | *.ens-mac.rest.gti.myantivirus.com
0 | | = | 0 | 2 | *.ens.rest.gti.myantivirus.com
0 | | = | 0 | 3 | *.evilzone.net
0 | | = | 0 | 2 | *.mace.rest.gti.myantivirus.com
0 | | = | 0 | 0 | *.mydomain.com
0 | | = | 0 | 0 | *.nicedomain.com
0 | | = | 0 | 3 | *.reallynotnicezone.com
0 | | = | 0 | 2 | *.realprotect1.myantivirus.com
0 | | = | 0 | 1 | *.secretzone.mydomain.com
0 | | = | 0 | 1 | *.secretzone2.outside.com
0 | | = | 0 | 1 | *.specialzone.mydomain.com
0 | | = | 0 | 2 | *.tie.gti.myantivirus.com
0 | | = | 0 | 2 | *.tunnel.hips.trustedsource.org
0 | | = | 0 | 2 | *.tunnel.message.trustedsource.org
0 | | = | 0 | 2 | *.tunnel.web.trustedsource.org
0 | | = | 0 | 2 | *.tunnel.web.trustedsource.org
0 | | = | 0 | 3 | *.verybadzone.ws
...

```

Table 58.2. DNS Guardian lists entries columns

Configuration parameter	Description
Used	The number of times the entry matched a query sent to Guardian. This implies the use of lists to filter views queries. For more details, refer to the section Filtering Guardian Views Using Lists .
*	The flag * indicates that the entry matches any subdomain of the domain displayed in the column <i>Name</i> .
=	The flag = indicates that the entry matches exactly the domain displayed in the column <i>Name</i> .
List ID members	The ID of the lists where the entry is present, separated by a , (comma).

Configuration parameter	Description
Name	The name of the list entry matching the domain to filter.

The results are sorted, first by number of queries matching this entry, then by domain name. You can filter it to only display the content of the list(s) of your choice, as in the example below:

```
DNS Blast> show list list_id=1

Used | * | = | EOL | List ID members | Name
  0 | | = | 0 | | 1 | *.specialzone.mydomain.com
  0 | | = | 0 | | 1 | *.secretzone.mydomain.com
  0 | | = | 0 | | 1 | *.secretzone2.outside.com
```

Keep in mind that you can access and set the list parameters in the GUI, on the properties page of the DNS server. For more details, refer to the chapter [Managing Guardian Configuration](#).

Table 58.3. Guardian lists configuration parameters

Configuration parameter	Description
List<0-7>.name	The name of the list. For more details, refer to the section Renaming a Guardian List .
List<0-7>.request xfer	The server parameters for the command <code>dig</code> used to automatically update the list. For more details, refer to the section Editing a Guardian List Content .
List<0-7>.save	The parameter that sets if a list is saved or cleared when the service DNS Guardian restarts. For more details, refer to the section Saving a Guardian List .
List<0-7>.zone name	The name of the zone from which the content is retrieved to automatically update the list. For more details, refer to the section Editing a Guardian List Content .

To display Guardian list entries

1. Open a shell session on your appliance using `root` credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To display the lists entries, use the following command:

```
show list
```

4. To display the entries of a specific list, use the following command:

```
show list list_id=<0-7>
```

Renaming a Guardian List

By default, the 8 Guardian lists are named `list0` to `list7`. The lists name and parameters, but not their content, are visible when you display Guardian configuration. For more details, refer to the section [Displaying Guardian Lists](#).

At any time, you can edit a list name.

To rename a Guardian list

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on . The properties page opens.

3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To edit the name of a list, in the field **List<0-7>.name**, specify the new name.
7. Click on **OK** to complete the operation. The wizard closes.

Editing a Guardian List Content

By default, there are 8 empty Guardian lists. You can:

- Manually fill in a list, one entry after another, via CLI.
- Automatically synchronize a list via a zone transfer from the GUI. The zone is located on a 3rd party DNS server (e.g. SURBL). The list is refreshed every minute.

Note that, in the list entries, you can use the wildcard *** in place of subdomain names, as in **.verybadzone.ws*.

For more details on how to display the list content and parameters, refer to the section [Displaying Guardian Lists](#).

To edit a Guardian list content

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To manually create a list entry, indicate a name respecting the syntax *<FQDN>* or *<*.FQDN>* and use the following command:

```
create list_entry list_id=<0-7> <list-entry-name>
```

4. To manually delete a list entry, indicate its exact name and use the following command:

```
clear list list_id=<0-7> <list-entry-name>
```

To update a Guardian list from an external source

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To indicate the IP address of the DNS server to query and, if need be, additional parameters to use for the dig command such as *-y hmac:<name:tsig-key>*, in the field **List<0-7>.request xfer**, specify *@<server-ip-address> <additional-dig-parameters>*¹.

¹For more details on the command *dig*, refer to the ISC documentation, available at [ftp://ftp.isc.org/www/bind/arm95/man.dig.html](http://ftp.isc.org/www/bind/arm95/man.dig.html).

7. To indicate the name of the DNS zone, located on the server you configured above, from which you want to retrieve the content to update the list, in the field **List<0-7>.zone name**, specify the name of the zone.
8. Click on **[OK]** to complete the operation. The wizard closes.

Saving a Guardian List

By default, whenever a change is made to a Guardian list, the list is saved in a file. However, you can clear the list(s) of your choice every time the service DNS Guardian restarts.

To save or clear Guardian list when the service DNS Guardian restarts

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **[⋮]**. The properties page opens.
3. In the panel **Options**, click on **[EDIT]**. The wizard **Options configuration** opens.
4. Click on **[NEXT]** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To indicate if the list of your choice must be saved every time DNS Guardian restarts, in the field **List<0-7>.save**, type in *1*.
7. In the field **List<0-7>.save**, indicate if the list should be saved or cleared every time DNS Guardian restarts:
 - To save the list, select *1*.
 - To clear the list, select *0*.

We strongly recommend leaving this parameter enabled with the default value *1*.

8. Click on **[OK]** to complete the operation. The wizard closes.

Resetting a Guardian List Counter

You can reset the number of times an entry matched a query sent to Guardian. For more details, refer to the section [Displaying Guardian Lists](#).

To reset a Guardian list counter

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:


```
/usr/local/nessy2/bin/blastcli
```
3. To reset all the counters of a specific list, use the following command:


```
reset list list_id=<0-7>
```
4. To reset only the counter of entries that have expired (*1*) or have not yet expired (*0*), use the following command:

```
reset list list_id=<0-7> expired=<0|1>
```

Clearing a Guardian List

You cannot delete a Guardian list but you can clear all the entries it contains or only those that have expired. For more details, refer to the section [Displaying Guardian Lists](#).

To clear Guardian lists

1. Open a shell session on your appliance using *root* credentials.
2. Connect to your local DNS Guardian using the command:

```
/usr/local/nessy2/bin/blastcli
```

3. To clear the content of a specific list, use the following command:

```
clear list list_id=<0-7>
```

4. To clear only the entries that have expired (1) or have not yet expired (0), use the following command:

```
clear list list_id=<0-7> expired=<0|1>
```

Managing Guardian Views

You can **configure the first 8 views of a Guardian server with specific traffic restrictions**.

You can add, order and manage views on a Guardian server. The position of a view is visible in the column *Order* on the page *All views*. For more details, refer to the chapters [Managing DNS Views](#) and [Configuring DNS Views](#).

Guardian allows to:

- Use Guardian lists to restrict domain-specific access for clients associated with a view. For more details, refer to the section [Filtering Guardian Views Using Lists](#).
- Log the queries to a domain belonging to a list, the answers and the lists matched. For more details, refer to the section [Enabling Querylog, Answerlog and List Log on Filtered Views](#).
- Set a transparent DNS proxy to hide the DNS server address from clients associated with a view. For more details, refer to the section [Setting a Transparent DNS Proxy for a DNS Guardian View](#).

Displaying Guardian Views Configuration

View parameters are available on the view and server properties page.

To display a view parameters

1. In the sidebar, go to **DNS > Views**. The page **All views** opens.
2. At the end of the line of the view of your choice, click on **ⓘ**. The properties page opens. The panel *Main properties* contains the parameters *Name*, *Match clients* and *Match destinations*.

For more details, refer to the chapters [Managing DNS Views](#) and [Configuring DNS Views](#).

Two Guardian view parameters are not visible on a view properties page. They are only available on the Guardian server properties page:

Table 58.4. Guardian view parameters available on the server properties page

Parameter	Description
View<0-7>.list id filter	The list filtering configuration for the view. It applies a traffic policy to specific list entries. For more details, refer to the section Filtering Guardian Views Using Lists .
View<0-7>.nat destination	The transparent DNS proxy configuration of the view which can be applied on a specific network. For more details, refer to the section Setting a Transparent DNS Proxy for a DNS Guardian View .

For more details regarding parameter configuration at server level, refer to the section [Editing Guardian Configuration](#).

Filtering Guardian Views Using Lists

When setting a DNS view:

- You can *allow* or *deny* access to the clients of your choice, by specifying ACLs, TSIG keys, IP and network addresses for the parameter *Match clients*.
- You can *allow* or *deny* access to the destination of your choice, also by specifying ACLs, TSIG keys, IP and network addresses using the parameter *Match destination*.

For more details, refer to the chapters [Managing DNS Views](#) and [Configuring DNS Views](#).

Since domain names can themselves be associated with many IP addresses, Guardian lists allow you to extend the *match-destination* configuration of a view to apply an advanced traffic policy depending on what domain was queried and who sent the query.

For each of the 8 first views on the server, you can:

1. Define which entries, from the Guardian list(s) of your choice, should be applied a traffic policy. For more details, refer to the section [Managing Guardian Lists](#).

Table 58.5. Guardian list filters for views

Filter	Description
any	The view policy applies to the entries present in any of the specified list(s).
all	The view policy only applies to the entries present in all of the specified lists. This implies the use of several lists.
none	The view policy applies to the entries present in none of the specified list(s).
default	The view policy applies to all the entries not impacted by another action. This implies the use of several actions. There is no need to specify a list for this filter.

2. Set one of the following Guardian view policies:

Table 58.6. Guardian view actions

Action	Description
nxdomain	The server sends an NXDOMAIN response to all the queries from the client.
nodata	The server sends a NODATA response to all the queries from the client.
redirect	The server redirects all the queries from the client. Both IPv4 and IPv6 redirection addresses can be set using, respectively, the parameters <i>List redirect a</i> and <i>List redirect aaa</i> . For more details, refer to the chapter Managing Guardian Configuration .
nocache	The server proceeds directly to recursion and does not answer with the cached records. The query is counted as <i>C-miss</i> in the cache statistics of the view. For more details, refer to the section Displaying Guardian Cache Content .

Action	Description
passthru	If the queried domain is present in the cache, the server replies with the cached answer. Otherwise, the server proceeds to recursion, send the answer to the client and caches the record for future use. The query is counted as <i>C-miss</i> in the cache statistics of the view. For more details, refer to the section Displaying Guardian Cache Content .
drop	The server drops all the queries from the client.

This policy applies to clients present in the *match-client* configuration of the view and querying an entry present in the specified list(s). When a client queries a domain that is not present in the list(s), Guardian applies the default policy *passthru*.

- Set as many view policies as needed, each configured with the list(s) of your choice, as described above.

To filter a Guardian view using lists

- In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
- At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
- In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
- Click on **NEXT** until you reach the last page of the wizard.
- In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
- To filter a Guardian view, in the field **View<0-7>.list id filter**, indicate one of the filters described in the table [Guardian list filters for views](#), one or more list IDs separated by a space and one of the policies described in the table [Guardian view policies](#), as in the following example:

```
<any/all/none> <space-separated-list-ids> <action>
```

- To configure different actions for the same view, separate them by a comma ",". Each action must be preceded by its own parameters, as in the following example:

```
<any/all/none> <space-separated-list-ids> <first-action>, <any/all/none> <space-separated-list-ids> <second-action>
```

- At the end of the command, you can indicate an action for all the entries not impacted by another action. There is no need to specify a list for this filter, as in the following example:

```
<any/all/none> <space-separated-list-ids> <first-action>, default <second-action>
```

- To remove all list restrictions from a view, leave the field **View<0-7>.list id filter** empty.
- Click on **OK** to complete the operation. The wizard closes.

Enabling Querylog, Answerlog and List Log on Filtered Views

The options *Querylog*, *Answerlog* and *List log* allow you to log the queries to a domain belonging to a list, the name of this list and/or the answers sent to the client. To do so:

- You must set Guardian parameters *Querylog*, *Answerlog* and/or *List log* to 2. For more details, refer to the chapter [Managing Guardian Configuration](#).
- You must type in the options *querylog*, *answerlog* and/or *list log* when you set the parameter *View<0-7>.list id filter*.

For instance, you can decide to enable the querylog, answerlog and list log for all the entries of the lists *mylist1* and *mylist2*, with *mylist1* containing wikipedia.com and reddit.com and *mylist2* containing reddit.com. If a client requests these two domains, the following information is available in the logs *named* on the page *Syslog*:

```
May 31 15:17:32 solid named[63380]: client 10.0.252.11#38824 (wikipedia.com): answer: wikipedia.com
IN A (10.0.81.4) -> SERVFAIL
May 31 15:17:32 solid named[63380]: client 10.0.252.11#38824: query: wikipedia.com IN A (10.0.81.4)
May 31 15:17:32 solid named[63380]: List Matched 10.0.81.4#38824: query: wikipedia.com IN A
(10.0.81.4){mylist1}
May 31 15:16:46 solid named[63380]: client 10.0.252.11#26822 (reddit.com): answer: reddit.com IN
A (10.0.81.4) -> SERVFAIL
May 31 15:16:46 solid named[63380]: client 10.0.252.11#26822: query: reddit.com IN A (10.0.81.4)
May 31 15:16:46 solid named[63380]: List Matched 10.0.81.4#26822: query: reddit.com IN A
(10.0.81.4){mylist1,mylist2}
```

For more details on how to display the logs, refer to the section [Syslog](#).

To enable Querylog, Answerlog and/or List log on a filtered view

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To enable the options Querylog, Answerlog and/or List log on the filtered view, in the field **View<0-7>.list id filter**, add *+querylog*, *+answerlog* and/or *+listlog* as in the following examples.

To enable the option querylog, type in:

```
<any/all/none> <space-separated-list-ids> <action> +querylog
```

To enable all three options, querylog, answerlog and list log, type in:

```
<any/all/none> <space-separated-list-ids> <action> +querylog+listlog+answerlog
```

7. Click on **OK** to complete the operation. The wizard closes.

Setting a Transparent DNS Proxy for a Guardian View

On a specific network, the administrator can root all the DNS traffic to a DNS server that answers all the queries at destination of port 53, whatever the destination address is. If a client connected to said network is configured with another DNS server address, it is still answered by Guardian with no mention of the routing. You can set one network address per view.

Note that this feature does not use NAT.

To set a transparent DNS proxy for a Guardian view

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⚙**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.

4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display option(s)**, select *all*. The list of all Guardian parameters is displayed.
6. To set a transparent DNS proxy for a view, in the field **View<0-7>.nat destination**, specify the network address and prefix of your choice as follows: *<network address>/<prefix>*.
7. To disable a transparent DNS proxy for a view, leave the field **View<0-7>.nat destination** empty.
8. Click on **OK** to complete the operation. The wizard closes.

Managing Guardian Policies

From the page *All policies*, you can manage and configure Guardian policies. A Guardian policy is a container for triggers and allows you to deploy a set of triggers on a Guardian server.

Two read-only policies are available by default in SOLIDserver: *default* and *empty*.

Browsing Guardian Policies

To display the list of Guardian policies

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. To display or hide policy deployments, on the right-end side of the menu, click on . The page refreshes.

In the column Name, the icon  precedes every policy.

If there are deployed policies, they are preceded by the icon  and listed under the parent policy name as many times as there are Guardian servers associated with the policy.

Two policies are available by default and read-only. You can [add policies](#) from scratch as described below or [duplicate already existing policies](#). For more details, refer to the section [Duplicating Guardian Policies](#).

To display a Guardian policy properties page

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. At the end of the line of the policy of your choice, click on . The properties page opens.

There are 5 columns on the page *All policies* that you can sort and filter. By default, all the columns are displayed on the page and you cannot change their order.

Table 58.7. Available columns on the page All policies

Column	Description
Name	The name of the policy.
Description	The description of the policy.
Guardian server	For a policy, the value is <i>N/A</i> . For a policy deployment the value is the name of the Guardian server the policy is deployed on.
Read only	Only the policies available by default are read-only. Their edition is possible but limited.
Status	Status of the policy.

Understanding the Policy Statuses

The column *Status* provides information regarding the policies you manage.

Table 58.8. Policy Statuses

Status	Description
 <i>OK</i>	The policy is operational.
 <i>Delayed create</i>	The policy is being created on the associated server.
 <i>Delayed delete</i>	The policy is being deleted on the associated server.

Adding and Deploying Guardian Policies

From the page *All policies*, you can add as many policies as you want. Before adding a policy, keep in mind that:

- A policy must have a unique name.
- When you deploy a policy on a Guardian server from the GUI, it overwrites the configuration on the server. That is to say, the triggers that were added via CLI are removed.
- No more than one policy can be deployed on a Guardian server and only servers in version 7.1 or higher are supported.
- You can deploy a policy at any time on a server. You can add a policy and its triggers before deploying it on one or several servers. For more details on how to add triggers, refer to the section [Managing Triggers](#).

To add a policy

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. In the menu, click on **+ Add**. The wizard **Add a policy** opens.
3. In the field **Name**, specify the name of the policy.
4. In the field **Description**, specify the description of the policy.
5. In the list **Available Guardian servers**, you can select a server and click on . The server is moved to the list **Selected Guardian servers**.

Only Guardian enabled physical servers are listed, whether they are managed by a smart architecture or not.

To remove a server from the **Selected Guardian servers**, select it and click on . The server is moved back to the list **Available Guardian servers**.

6. Click on  to complete the operation. The policy is listed.

If you deployed the policy on one or more Guardian servers, click on . Several lines appear under the policy itself, there is a line for each of the server(s) the policy is deployed on.

Editing Guardian Policies

At any time, you can edit policies. Before editing a policy, note that:

- You cannot rename policies.
- You cannot edit  policy deployments. You must edit the policy itself.

- Editing a policy means:
 - Editing its definition.
 - Associating it with extra Guardian servers.
 - Dissociating it from one or more Guardian servers.
- The edition of read-only policies is limited to:
 - Associating it with extra Guardian servers.
 - Dissociating it from one or more Guardian servers.
- No more than one policy can be deployed on a Guardian server and only servers in version 7.1 or higher are supported.

To edit a policy

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. Right-click on the name of the policy that you want to edit. The contextual menu appears.
3. Click on  **Edit**. The wizard **Edit a policy** opens.
4. Edit the policy **Description** according to your needs. If you are editing a read-only policy, this field is grayed out.
5. Edit the list **Selected Guardian servers** according to your needs. Only Guardian enabled physical servers are listed, whether they are managed by a smart architecture or not.

Select a server in the list *Available Guardian servers* and click on  to add it the list *Selected Guardian servers*. Select a server in the list *Selected Guardian servers* and click on  to remove it move it back to the list *Available Guardian servers*.

If you remove a server from a policy, this server is not be available for a new association to a different policy (or to this same server) as long as the policy deployment has the status *Delayed delete*.

6. Click on  to complete the operation. The page refreshes.

Duplicating Guardian Policies

You can make a copy of an existing policy. The new policy will contain the same triggers as the duplicated one. In this new policy, the positions of the triggers are reset and start from 0, and are incremented for each trigger copied. For more details, refer to the section [Managing Triggers](#). Note that:

- You cannot duplicate  policy deployments. You must duplicate the policy itself.
- The association between the policy and a Guardian server is not duplicated.
- The read-only status of the new policy is set to *No*.

To duplicate a policy

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. In the list, select the policy that you want to duplicate.
3. In the menu, select  **Edit > Duplicate**. The wizard **Duplicate a policy** opens.
4. In the field **Policy**, name the new policy.

5. Click on **OK** to complete the operation. The new policy is listed.

Deleting Guardian Policies

At any time, you can delete a policy. Deleting a policy also deletes the triggers it contains. Note that:

- You cannot delete read-only policies.
- You cannot delete policies associated with a Guardian server. You must dissociate the server first. For more details, refer to the section [Editing Guardian Policies](#).
- You cannot delete  policy deployments. Once the policy is dissociated from the server, the dedicated deployment line is no longer listed. If a server is not part of your resources, you may not see some policy deployments that prevent from deleting policies.

To delete a policy

1. In the sidebar, go to  **Guardian > Policies**. The page **All policies** opens.
2. In the list, select the policy that you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The policy is no longer listed.

Managing Triggers

From the page *All triggers*, you can manage and configure triggers.

DNS Guardian client statistics can be used to protect the server by **targeting suspicious clients and restricting their traffic**. From version 7.1 of SOLIDserver, you can manage triggers from the GUI, on the page *All triggers*.

To do so, you can add triggers that:

1. Are armed when a client traffic reaches a certain traffic threshold.
2. Apply an action on the client traffic for a specific period of time.
3. Are disarmed when the action is over.

Once a client traffic reaches a threshold, the trigger is armed for the whole specified period even if the client traffic drops down below said threshold. If the client reaches the traffic threshold again when the trigger is already armed, the action duration is renewed but there is no mention of the rearming in the logs *named* on the page *Syslog*. For more details, refer to the section [Syslog](#).



Figure 58.2. Triggers

Note that, for the triggers to be available, DNS Guardian Protection must be enabled. For more details, refer to the section [Enabling Guardian Protection](#).

Keep in mind that you can monitor the **analytics of all trigger armings** as well as the **graph statistics of 8 triggers**. For more details, refer to the section [Monitoring Guardian from the GUI](#).

Limitations

- **Any change made in the GUI overwrites the changes made via CLI.** It is still possible to add triggers via CLI. However, the modifications made via CLI are not displayed in the GUI. For instance, if you add a trigger via CLI, it is not displayed on the page *All triggers*. If you modify the policy deployed on this server or one of its triggers in the GUI, the modifications made in the CLI are erased and replaced by the data available in the GUI.
- If you restore a backup on a management appliance, the modifications, made in the module *Guardian* between the time the backup was made and the time the backup is restored, are still applied on the server but are not displayed in the GUI. If after the backup restoration, you make a modification in the module *Guardian*, the configuration currently displayed in the GUI is deployed on the servers.

Browsing Triggers

To display the list of triggers

1. In the sidebar, go to **Guardian** > **Triggers**. The page **All triggers** opens.
2. To display or hide policy deployments, on the right-end side of the menu, click on . For more details about policy deployments, refer to the section [Browsing Policies](#).

By default, there are five triggers available in the policy *default*. You cannot modify these triggers as the policy *default* is in read-only.

To display a trigger properties page

1. In the sidebar, go to **Guardian** > **Triggers**. The page **All triggers** opens.

- At the end of the line of the trigger of your choice, click on . The properties page opens.

Customizing the Display on the Page All Triggers

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

There are 9 columns on the page *All triggers* that you can sort and filter. By default, all the columns are displayed on the page.

Table 58.9. Available columns on the page All triggers

Column	Description
Name	The name of the trigger.
Action	The action that is launched when the trigger is armed. For more details, refer to the table DNS Guardian trigger actions .
Duration	The duration (in seconds) of the action.
Options	Shows if querylog and/or answerlog are enabled. For more details, refer to the section Enabling Querylog and Answerlog on Triggers .
Rule definition	The definition of the rule that will trigger the action.
Policy	The name of the policy or policy deployment the trigger belongs to.
Guardian server	The name of the server on which the policy the trigger belongs to is deployed. For a policy line, the value is <i>N/A</i> .
Status	Status of the trigger. For more details, refer to the table Trigger statuses .

Understanding the Trigger Statuses

The column *Status* provides information regarding the triggers you manage.

Table 58.10. Trigger statuses

Status	Description
 <i>OK</i>	The trigger is managed and operational.
 <i>Delayed create</i>	The trigger is being created on the associated server.
 <i>Delayed delete</i>	The trigger is being deleted on the associated server.
 <i>Unmanaged</i>	The trigger is unmanaged.

Adding Triggers

You can add as many triggers as needed. Triggers contain:

- One of the actions illustrated below.
- The rule that defines when the action is launched. The trigger rule syntax is based on the Reverse Polish Notation (RPN) which is parenthesis-free and where each operator follows all of its operands. The rule uses one or more metrics from the Guardian clients metrics described in the table [Trigger metrics](#).
- The metric operators +, -, *, /.
- The threshold operators =, !=, >=, <=, <, >.

With triggers, you can launch five actions in order to protect your servers:

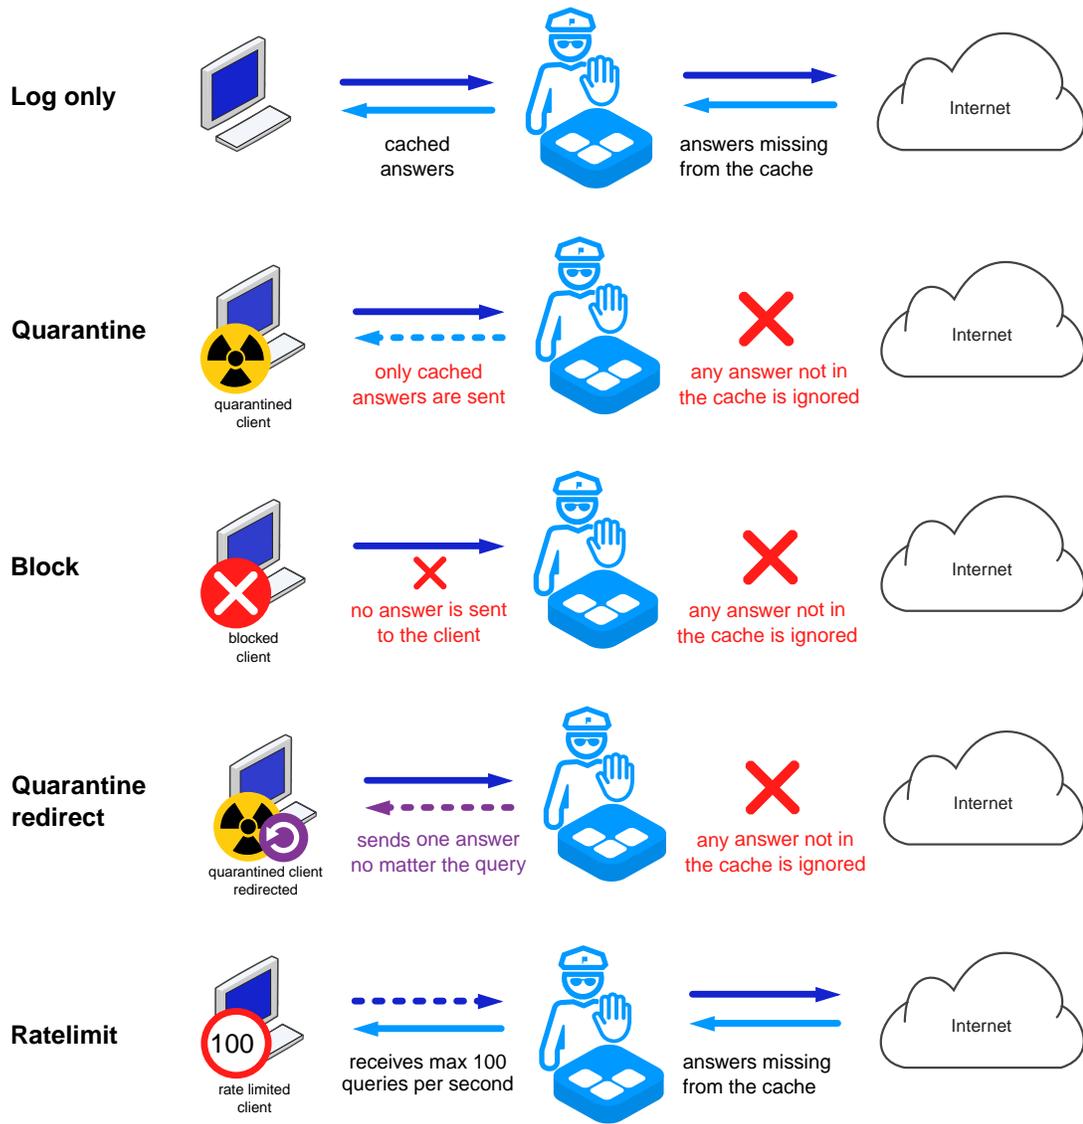


Figure 58.3. Guardian actions

In the table below are described the metrics you can use in the triggers:

Table 58.11. Trigger metrics

Metric	Description
<0-999999>	A number between 0 and 999999.
query	Number of queries received.
auth	Number of authoritative queries received.
invalid_query	Number of invalid queries received.
cache_hit	Number of cache hits.
cache_miss	Number of cache misses.
cache_hit_quarantine	Number of cache hits in quarantine mode.
cache_miss_quarantine	Number of cache misses in quarantine mode.

Metric	Description
recurs_time	Total time in recursion in milliseconds.
auth_query_size	Size, in bytes, of authoritative queries.
auth_answer_size	Size, in bytes, of authoritative answers.
hit_query_size	Size, in bytes, of hit queries.
hit_answer_size	Size, in bytes, of hit answers.
miss_query_size	Size, in bytes, of missed queries.
miss_query_size_not_exist	Size, in bytes, of missed queries that have never been cached.
miss_answer_size	Size, in bytes, of answers to missed queries.
miss_answer_size_not_exist	Size, in bytes, of answers to missed queries that have never been cached.
answer_noerror	Number of answers returning a NOERROR rcode.
answer_formerr	Number of answers returning a FORMERR rcode.
answer_servfail	Number of answers returning a SERVFAIL rcode.
answer_nxdomain	Number of answers returning a NXDOMAIN rcode.
answer_notimp	Number of answers returning a NOTIMP rcode.
answer_refused	Number of answers returning a REFUSED rcode.
auth_answer_noerror	Number of authoritative answers returning a NOERROR rcode.
auth_answer_formerr	Number of authoritative answers returning a FORMERR rcode.
auth_answer_servfail	Number of authoritative answers returning a SERVFAIL rcode.
auth_answer_nxdomain	Number of authoritative answers returning a NXDOMAIN rcode.
auth_answer_notimp	Number of authoritative answers returning a NOTIMP rcode.
auth_answer_refused	Number of authoritative answers returning a REFUSED rcode.
hit_answer_noerror	Number of answers (hits) returning a NOERROR rcode.
hit_answer_formerr	Number of answers (hits) returning a FORMERR rcode.
hit_answer_servfail	Number of answers (hits) returning a SERVFAIL rcode.
hit_answer_nxdomain	Number of answers (hits) returning a NXDOMAIN rcode.
hit_answer_notimp	Number of answers (hits) returning a NOTIMP rcode.
hit_answer_refused	Number of answers (hits) returning a REFUSED rcode.
miss_answer_noerror	Number of answers (misses) returning a NOERROR rcode.
miss_answer_formerr	Number of answers (misses) returning a FORMERR rcode.
miss_answer_servfail	Number of answers (misses) returning a SERVFAIL rcode.
miss_answer_nxdomain	Number of answers (misses) returning a NXDOMAIN rcode.
miss_answer_notimp	Number of answers (misses) returning a NOTIMP rcode.
miss_answer_refused	Number of answers (misses) returning a REFUSED rcode.
rpz_hit	Number of RPZ hits.

For instance, you can add the trigger Trigger-B:

1. **That applies the traffic action Block during 3600 seconds** to the client that arms the trigger,
2. That is armed when the client generates more than 100 queries during 10 consecutive seconds.

To do so, select the action *Block* for a duration of *3600s* and type in the rule *query@10 100 >*.

All the events that armed (*ARMING*) or disarmed (*DISARMING*) the trigger are detailed in the logs *named* on the page *Syslog*, along with the client's IP address and the related action policy. For more details on how to display the logs, refer to the section [Syslog](#).

The arming and disarming of the trigger are logged in the logs *named* with the tag *Trigger-B*:

```
01/03/2018 18:13:18 solid named[12962]: ARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-B)
01/03/2018 18:13:18 solid named[12962]: DISARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-B)
```

Before adding a trigger, note that:

- You cannot add a trigger in a read-only policy.
- You can use Guardian lists to limit a metric value to the domains of your choice. For more details, refer to the section [Adding Triggers Relying on Lists Metrics](#).
- You can use the operators & and / to set different metric thresholds for the same trigger. For more details, refer to the section [Adding Triggers Armed by Several Thresholds](#).
- You can use the operators % and *push_param* to indicate, in the logs *named*, the threshold that has been reached. For more details, refer to the section [Adding Tagged Triggers](#).

To add a trigger

1. In the sidebar, go to  **Guardian > Triggers**. The page **All triggers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Guardian policy**, select the policy in which you want to add the trigger. Only the non read-only policies are available.
4. Click on **NEXT**. The page **Add a trigger** opens.
5. The field **Position** is in read-only. The value is retrieved from the available positions.
6. In the field **Name**, specify the name of the trigger.
7. In the drop-down list **Action**, select the action that you want to launch when the trigger is armed. For more details, refer to the table below.

Table 58.12. Trigger actions

Action	Description
Log only	The arming and disarming of the trigger is logged in the logs <i>named</i> on the page <i>Syslog</i> . This is the default value.
Quarantine	The arming and disarming of the trigger is logged and the client is placed in quarantine to protect the local recursive servers: cached queries are allowed but recursion is blocked.
Block	The arming and disarming of the trigger is logged and DNS resolution is blocked. Be careful using this action since blocking a client's address might ensue in blocking all other clients using it to connect the network , for instance, in the case of private hosts on the same NAT device or router.
Quarantine redirect	The arming and disarming of the trigger is logged and the address set for the option <i>Quarantine redirect A</i> or <i>Quarantine redirect AAAA</i> is returned for queries not cached. For more details, refer to the chapter Managing Guardian Configuration .
Ratelimit	The arming and disarming of the trigger is logged and the client traffic is capped to 100 queries per second. This value is set by default and cannot be changed.

8. In the field **Duration (in seconds)**, specify how long the action should last.
9. In the section **Action options**, tick the boxes **Querylog** and/or **Answerlog**. For more details, refer to the section [Enabling Querylog and Answerlog on Triggers](#).
10. In the field **Rule definition**, type in the definition of the rule in Reversed Polish Notation, using at least one of one of the metrics described in the table [Trigger metrics](#) as in the following example:

`<metric>@<number-of-seconds> <threshold-value> <threshold-operator>`

You can also add a trigger using the result of an operation combining two or more metrics using a metric operator, as in the following example:

`<first-metric>@<number-of-seconds> <second-metric>@<number-of-seconds> <metric-operator> <threshold-value> <threshold-operator>`

The syntax above can also be combined with Guardian lists. For more details, refer to the section [Adding Triggers Relying on Lists Metrics](#).

11. To activate the trigger, tick the box **Enabled**. If you do not, it is never armed even if its threshold is reached.

Note that you can manage this option on one or more triggers at once. For more details, refer to the section [Managing or Unmanaging Triggers](#).

12. Click on to complete the operation. The report works for a while before closing. The list is updated.

If you added the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

Adding Triggers Relying on List Metrics

When setting the threshold of a trigger, you can limit the metric of your choice to a list of domains. This is useful to avoid false positives by excluding, from the threshold, queries towards the legitimate domains of your choice.

For instance, you can add a trigger:

1. That applies the traffic action *Ratelimit* during 3600 seconds to the client that arms the trigger.
2. That is armed when the client generates more than 100 queries not present in the cache during 10 consecutive seconds, **from which are deducted those made towards the domains present in "list1"**.

To do so, select the action *Ratelimit* for a duration of 3600s and type in the rule `cache_miss@10 list1_cache_miss@10 - 100 >`.

The list metrics available are described in the table below:

Table 58.13. Guardian list metrics

List metric	Description
list<0-7>_cache_hit	Cumulated number of cache hits for all the domains present in the list <0-7>.
list<0-7>_cache_miss	Cumulated number of cache misses for all the domains present in the list <0-7>.
list<0-7>_hit_answer_size	Cumulated size of hit answers for all the domains present in the list <0-7>.
list<0-7>_hit_query_size	Cumulated size of hit queries for all the domains present in the list <0-7>.
list<0-7>_miss_answer_size	Cumulated size of missed answers for all the domains present in the list <0-7>.
list<0-7>_miss_query_size	Cumulated size of missed answers for all the domains present in the list <0-7>.

Before adding a trigger, note that:

- You cannot add a trigger in a read-only policy.

- You can use the operators `&` and `/` to set different metric thresholds for the same trigger. For more details, refer to the section [Adding Triggers Armed by Several Thresholds](#).
- You can use the operators `%` and `push_param` to indicate, in the logs *named*, the threshold that has been reached. For more details, refer to the section [Adding Tagged Triggers](#).

To add a trigger relying on lists metrics

1. In the sidebar, go to  **Guardian** > **Triggers**. The page **All triggers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Guardian policy**, select the policy in which you want to add the trigger. Only the non read-only policies are available.
4. Click on **NEXT**. The page **Add a trigger** opens.
5. The field **Position** is in read-only. The value is retrieved from the available positions.
6. In the field **Name**, specify the name of the trigger.
7. In the drop-down list **Action**, select the action that you want to launch when the trigger is armed.
8. In the field **Duration (in seconds)**, specify how long the action should last.
9. In the section **Action options**, tick the boxes **Querylog** and/or **Answerlog**. For more details, refer to the section [Enabling Querylog and Answerlog on Triggers](#).
10. In the field **Rule definition**, type in the definition of the rule in Reversed Polish Notation using at least one or more metrics and one or more list IDs. For more details, refer to the sections [Adding Triggers](#) and [Managing Guardian Lists](#).

The following rule describes how to add a trigger with one metric and one list:

```
<metric> @<number-of-seconds> <list-name>_<metric> @<number-of-seconds> <metric-operator> <threshold-value> <threshold-operator>
```

11. To activate the trigger, tick the box **Enabled**. If you do not, it is never armed even if its threshold is reached.

Note that you can manage this option on one or more triggers at once. For more details, refer to the section [Managing or Unmanaging Triggers](#).

12. Click on **OK** to complete the operation. The report works for a while before closing. The list is updated.

If you added the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

Adding Triggers Armed by Several Thresholds

When configuring a trigger, you can specify if it should be armed when all or only some of several thresholds are reached.

For instance, you can add the trigger Trigger-S:

1. That applies the traffic action *Quarantine* during 3600 seconds to the client that arms the trigger.
2. That is armed when the client generates **more than 100 queries during 10 consecutive seconds OR more than 50 queries missing the cache during 30 consecutive seconds**.

To do so, select the action *Quarantine* for a duration of 3600s and type in the rule *query@10 100 > cache_miss@30 50 > |*.

Before adding a trigger, note that:

- You cannot add a trigger in a read-only policy.
- You can use Guardian lists to limit a metric value to the domains of your choice. For more details, refer to the section [Adding Triggers Relying on Lists Metrics](#).
- You can use the operators % and *push_param* to indicate, in the logs *named*, the threshold that has been reached. For more details, refer to the section [Adding Tagged Triggers](#).

To add a trigger armed by several thresholds

1. In the sidebar, go to  **Guardian > Triggers**. The page **All triggers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Guardian policy**, select the policy in which you want to add the trigger. Only the non read-only policies are available.
4. Click on **NEXT**. The page **Add a trigger** opens.
5. The field **Position** is in read-only. The value is retrieved from the available positions.
6. In the field **Name**, specify the name of the trigger.
7. In the drop-down list **Action**, select the action that you want to launch when the trigger is armed.
8. In the field **Duration (in seconds)**, specify how long the action should last.
9. In the section **Action options**, tick the boxes **Querylog** and/or **Answerlog**. For more details, refer to the section [Enabling Querylog and Answerlog on Triggers](#).
10. In the field **Rule definition**, type in the definition of the rule in Reversed Polish Notation using several metrics. For more details, refer to the section [Adding Triggers](#).
 1. Use the structure `<metric>@<number-of-seconds> <metric-threshold> <threshold-value> <threshold-operator>` as many times as needed.
 2. Place the operator `|` and/or `&` after all the operands as many times as needed.

The following rule describes how to add a trigger with two thresholds:

```
<first-metric>@<number-of-seconds> <first-metric-threshold> <first-threshold-value> <first-threshold-operator> <second-metric>@<number-of-seconds> <second-metric-threshold> <second-threshold-value> <second-threshold-operator> <()or(&>
```

11. To activate the trigger, tick the box **Enabled**. If you do not, it is never armed even if its threshold is reached.

Note that you can manage this option on one or more triggers at once. For more details, refer to the section [Managing or Unmanaging Triggers](#).

12. Click on **OK** to complete the operation. The report works for a while before closing. The list is updated.

If you added the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

The arming and disarming of the trigger is logged in the logs *named* on the page *Syslog* but does not indicate which threshold was met. To display this information, you can use trigger tags. For more details, refer to the section [Adding Tagged Triggers](#).

Adding Tagged Triggers

You can display, in the logs *named* on the page *Syslog*, the metric that armed a trigger as well as the threshold that has been reached.

Each metric you want to display in the tag is given a numerical position, from left to right, starting from 0.

For instance, you can add a trigger:

1. Called "**Trigger-T1 Queries %0**" where the operator **%0** is replaced in the logs by the **first tagged metric**.
2. That is enabled and applies the traffic policy *Block* to the client during 3600 seconds.
3. That is armed when the client generates **more than 100 queries during 10 consecutive seconds, which is the first metric preceded by the operator push_param**.

To do so, name your trigger *Trigger-T1 Queries %0*, select the action *Block* for a duration of 3600 seconds and type in the rule *query@10 push_param 100 >=*.

The arming and disarming of the trigger is logged in the logs *named* with the tag **Trigger-T1 Queries <value>**:

```
01/03/2018 18:13:18      solid named[12962]: ARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-T1
Queries 100)
01/03/2018 18:13:18      solid named[12962]: DISARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-T1
Queries 0)
```

For more details on how to display the logs, refer to the section [Syslog](#).

Tags are the most useful when the trigger has several metrics, as in **Trigger-T2**:

1. Called "**Trigger-T2 Queries %0 Cache-Miss %1**" where the operator **%0** is replaced in the logs by the **first tagged metric** and **%1** by the **second tagged metric**.
2. Which is enabled and applies the traffic policy *Block* to the client during 3600 seconds.
3. That is armed when the client generates **more than 100 queries during 10 consecutive seconds, which is the first metric** preceded by the operator **push_param** or when the client generates **more than 50 queries missing the cache in 30 consecutive seconds, which is the second metric** preceded by the operator **push_param**.

To do so, name your trigger *Trigger-T2 Queries %0 Cache-Miss %1*, select the action *Block* for a duration of 3600 seconds and type in the rule *query@10 push_param 100 >= cache_miss@30 push_param 50 >= |*.

The arming and disarming of the trigger is logged in the logs *named* with the tag **Trigger-T2 Queries <value> Cache-Miss <value>**.

```
01/03/2018 18:13:18      solid named[12962]: ARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-T2
Queries 0 Cache-Miss 50)
01/03/2018 18:13:18      solid named[12962]: DISARMING trigger on 142.12.0.101 (action:BLOCK) (Trigger-T1
Queries 0 Cache-Miss 0)
```

The result above indicates that the threshold of the second metric, tagged as *Cache-Miss*, was met.

Before adding a trigger, note that:

- You cannot add a trigger in a read-only policy.
- You can use Guardian lists to limit a metric value to the domains of your choice. For more details, refer to the section [Adding Triggers Relying on Lists Metrics](#).
- You can use the operators & and / to set different metric thresholds for the same trigger. For more details, refer to the section [Adding Triggers Armed by Several Thresholds](#).

To add a tagged Guardian trigger

1. In the sidebar, go to  **Guardian** > **Triggers**. The page **All triggers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Guardian policy**, select the policy in which you want to add the trigger. Only the non read-only policies are available.
4. Click on **NEXT**. The page **Add a trigger** opens.
5. The field **Position** is in read-only. The value is retrieved from the available positions.
6. In the field **Name**, specify the name and each *<metric name>* followed by *%<position-number>*:

<trigger-name> <first-metric-name> %0 <second-metric-name> %1
7. In the drop-down list **Action**, select the action that you want to launch when the trigger is armed.
8. In the field **Duration (in seconds)**, specify how long the action should last.
9. In the section **Action options**, tick the boxes **Querylog** and/or **Answerlog**. For more details, refer to the section [Enabling Querylog and Answerlog on Triggers](#).
10. In the field **Rule definition**, type in the definition of the rule in Reversed Polish Notation. Before each metric threshold, add the *push_param*:

<metric>@<number-of-seconds> push_param <first-threshold> <metric>@<number-of-seconds> push_param <second-threshold>|
11. To activate the trigger, tick the box **Enabled**. If you do not, it is never armed even if its threshold is reached.

Note that you can manage this option on one or more triggers at once. For more details, refer to the section [Managing or Unmanaging Triggers](#).
12. Click on **OK** to complete the operation. The report works for a while before closing. The list is updated.

If you added the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

Enabling Querylog and Answerlog on Triggers

By default, Guardian logs every time a trigger is armed (*arming*) and disarmed (*disarming*). For more details, refer to the section [Adding Triggers](#).

In addition, you can customize a trigger to also log the queries and/or answers of a client during the whole action duration. To do so:

- You must set Guardian parameters *querylog* and/or *answerlog* to 2. For more details, refer to the chapter [Managing Guardian Configuration](#).
- You must select the options *querylog* and/or *answerlog* when you add the trigger.

For instance, you can decide to enable the *querylog* and *answerlog* along with the action *Log only* during 60 seconds, for any client that **receives more than 10 SERVFAIL rcode answers during 100 consecutive seconds** from the server:

For instance, you can add the trigger *Trigger-Z*:

1. That applies the traffic action *Log only* to the client during 60 seconds **and logs its queries and answers at the same time**.
2. That is armed when the client generates more than 500 queries during 10 consecutive seconds.

To do so, select the action *Log only* for a duration of 60s as well as the actions *Querylog* and *Answerlog* and type in the rule *query@10 500 >*.

The arming and disarming of the trigger as well as the client's queries and/or answers are logged in the logs *named* with the tag *Trigger-Z*:

```
01/03/2018 15:17:37    solid named[12962]: DISARMING trigger on 10.0.252.11 (action:NONE) (Trigger-Z)
01/03/2018 15:17:32    solid named[12962]: client 10.0.252.11#38824 (wikipedia.com): answer:
wikipedia.com IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:17:32    solid named[12962]: client 10.0.252.11#38824: query: wikipedia.com IN A
(10.0.81.4)
01/03/2018 15:16:46    solid named[12962]: client 10.0.252.11#26822 (reddit.com): answer: reddit.com
IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:16:46    solid named[12962]: client 10.0.252.11#26822: query: reddit.com IN A (10.0.81.4)
01/03/2018 15:16:46    solid named[12962]: client 10.0.252.11#29848 (reddit.com): answer: reddit.com
IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:16:46    solid named[12962]: client 10.0.252.11#29848: query: reddit.com IN A (10.0.81.4)
01/03/2018 15:16:45    solid named[12962]: client 10.0.252.11#35200 (reddit.com): answer: reddit.com
IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:16:45    solid named[12962]: client 10.0.252.11#35200: query: reddit.com IN A (10.0.81.4)
01/03/2018 15:16:44    solid named[12962]: client 10.0.252.11#64691 (reddit.com): answer: reddit.com
IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:16:44    solid named[12962]: client 10.0.252.11#64691: query: reddit.com IN A (10.0.81.4)
01/03/2018 15:16:42    solid named[12962]: client 10.0.252.11#16781 (reddit.com): answer: reddit.com
IN A (10.0.81.4) -> SERVFAIL
01/03/2018 15:16:42    solid named[12962]: client 10.0.252.11#16781: query: reddit.com IN A (10.0.81.4)
01/03/2018 15:16:37    solid named[12962]: ARMING trigger on 10.0.252.11 (action:NONE) (Trigger-Z)
```

For more details on how to display the logs, refer to the section [Syslog](#).

To enable Querylog and Answerlog on Triggers

1. In the sidebar, go to  **Guardian > Triggers**. The page **All triggers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Guardian policy**, select the policy in which you want to add the trigger. Only the non read-only policies are available.
4. Click on **NEXT**. The page **Add a trigger** opens.
5. The field **Position** is in read-only. The value is retrieved from the available positions.
6. In the field **Name**, specify the name of the trigger.
7. In the drop-down list **Action**, select the action that you want to launch when the trigger is armed.
8. In the field **Duration (in seconds)**, specify how long the action should last.
9. In the section **Action options**, tick the boxes **Querylog** and/or **Answerlog**.

10. In the field **Rule definition**, type in the definition of the rule in Reversed Polish Notation, using at least one of one of the metrics described in the table [Trigger metrics](#).
11. To activate the trigger, tick the box **Enabled**. If you do not, it is never armed even if its threshold is reached.

Note that you can manage this option on one or more triggers at once. For more details, refer to the section [Managing or Unmanaging Triggers](#).

12. Click on to complete the operation. The report works for a while before closing. The list is updated.

If you added the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

Editing Triggers

At any time you can update a trigger. Before editing a trigger, note that:

- You cannot edit a trigger belonging to a read-only policy or a  policy deployment.
- The trigger edition consists of editing its main properties and its rule definition.
- Unticking the box *Enabled* unmanages a trigger. For more details, refer to the section [Managing or Unmanaging Triggers](#).

To edit a trigger

1. In the sidebar, go to  **Guardian > Triggers**. The page **All triggers** opens.
2. Right-click on the name of the trigger you want to edit. The contextual menu opens.
3. Click on **Edit**. The wizard **Edit a trigger** opens.
4. Edit the trigger field(s) **Name**, **Action**, **Duration (in seconds)**, **Action options**, **Rule definition** (in Reversed Polish Notation) and **Enabled** according to your needs.
5. Click on to complete the operation. The report works for a while before closing. The list is updated.

If you edited the trigger in a policy that is deployed on one or more Guardian servers, click on . Several lines appear under the trigger itself, there is a line for each of the server(s) the policy is deployed on.

Managing or Unmanaging Triggers

At any time you can stop managing one or more triggers, which means disabling them. Note that:

- An unmanaged trigger is not armed even if its threshold is reached.
- Editing a trigger to untick the box *Enabled* allows to individually unmanage the trigger.
- You cannot manage or unmanage a trigger if the *Guardian server* value is different from N/A.

To manage or unmanage a trigger

1. In the sidebar, go to  **Guardian > Triggers**. The page **All triggers** opens.
2. Tick the trigger(s) of your choice.

3. In the menu, select **Edit > Manage > Yes** or **No**. The **Items managements** wizard opens.
4. Click on **OK** to complete the operation. In the column **Status**, the trigger is marked **Unmanaged**.

Deleting Triggers

At any time, you can delete a trigger. Keep in mind that you cannot delete a trigger if the *Guardian server* value is different from N/A.

To delete a trigger

1. In the sidebar, go to **Guardian > Triggers**. The page **All triggers** opens.
2. Tick the trigger you want to delete.
3. In the menu, click on **Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The device is no longer listed.

Disabling Guardian Protection

Disabling Guardian protection allows to stop using the cache to answer client queries while still updating its content:

- Guardian still caches the local recursive server queries/answers but does not send any information to DNS clients. This can be useful to fill in the cache.
- The server receives and sends new entries to the other Guardian servers if you configured cache sharing. For more details, refer to the section [Sharing the Cache of Several Guardian servers](#).
- Guardian server ignores the Rescue mode configuration and can never switch to that mode. For more details, refer to the section [Managing DNS Guardian Rescue Mode](#).
- Guardian server ignores the triggers configured. For more details, refer to the section [Managing Triggers](#).

To disable Guardian protection

1. In the sidebar, go to **DNS > Servers**. The page **All servers** opens.
2. At the end of the line of the Guardian server of your choice, click on **⌵**. The properties page opens.
3. In the panel **Options**, click on **EDIT**. The wizard **Options configuration** opens.
4. Click on **NEXT** until you reach the last page of the wizard.
5. In the drop-down list **Display options(s)**, select which parameters to display: *all, using non-default values, using default values or different from smart*.
6. To disable DNS Guardian protection, in the field **Blast**, select *no (0)*.
7. Click on **OK** to complete the operation. The wizard closes.

Part XI. NetChange

NetChange allows to locate and monitor your network devices. You can import them using their IP address and the CDP, NDP and LLDP layer 2 discovery protocols to find more devices on the network.

Once imported, NetChange relies on the SNMP protocol to retrieve information in the MIB of each device. It provides monitoring for the routes, VLANs and ports the devices contain and the IP addresses of their interfaces as well as the devices connected to them. In addition, it allows to monitor the configuration file versioning on the devices that support it.

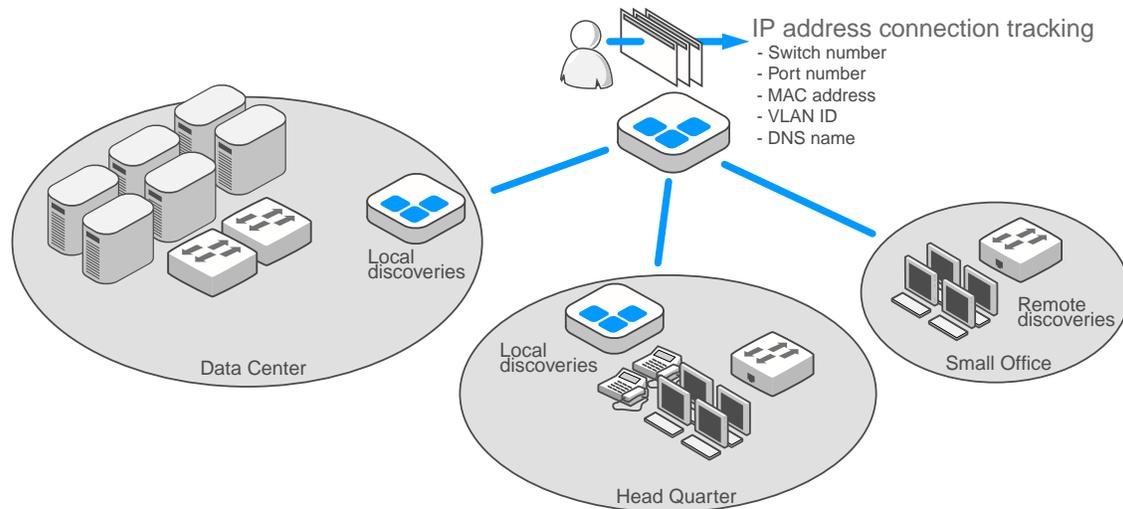


Figure 210. Example of a NetChange architecture

We recommend importing all possible devices to have a clear overview of your network. All supported devices are listed on the Knowledge Base in the category NetChange/IPLocator at https://kb.efficientip.com/index.php/Main_Page.

NetChange can include up to 3 levels of organization:

- **Network devices:** the highest level of hierarchy where you import all the devices you want to manage. They contain routes, VLANs, ports and/or configurations files. For more details, refer to the chapter [Managing Network Devices](#).
 - **Routes:** one of the second levels of hierarchy where you can view and partially manage the routing tables of your layer 3 devices. For more details, refer to the chapter [Managing Routes](#).
 - **VLANs:** one of the second levels of the hierarchy where you can manage the VLANs of your network devices. For more details, refer to the chapter [Managing VLANs](#).
 - **Ports:** one of the second levels of the hierarchy where you can manage the ports of your network devices. For more details, refer to the chapter [Managing Ports](#).
 - **Configurations:** one of the second levels of the hierarchy where you can manage the configuration file versioning of the network devices that support it. For more details, refer to the chapter [Managing Configuration Versioning](#).
 - **Addresses:** one of the second levels of the hierarchy where you can view the interface IP addresses of the imported network devices. For more details, refer to the chapter [Managing Addresses](#).
-

- **Discovered items:** the lowest level of the hierarchy where you can manage the devices connected to your network devices. The devices are identified through their MAC address and are connected to a network device via VLANs or ports. For more details, refer to the chapter [Managing Discovered Items](#).

The module also provides:

- **Statistics.** You can use gadgets to display device and/or port statistics. For more details, refer to the chapter [Managing Statistics](#).
- **Monitoring and Tuning.** You generate reports on devices, refresh them using CSV files, keep them up-to-date or customize their type. For more details, refer to the chapter [Monitoring and Tuning](#).

Note that the **depending on the license you chose**, you can either retrieve network devices information or partially configure your network devices and their content. There are two NetChange licenses available:

1. **NetChange-IPL**, a light version that provides basic management options of your network devices.
2. **NetChange**, the full license that allows advanced management of your Avaya, Cisco and HP network devices as it provides configuration options for VLANs and ports, 802.1X authentication, versioning...

Table 331. NetChange licenses differences

Option	NetChange-IPL	NetChange
Adding and listing Network devices	YES	YES
Discovering and listing routes	YES	YES
Listing VLANs	YES	YES
Listing ports	YES	YES
Listing IP addresses configured on the network device interfaces	YES	YES
Listing discovered items	YES	YES
Adding and deleting VLANs	NO	YES
Configuring the speed of a port	NO	YES
Configuring the duplex of a port	NO	YES
Configuring the VLAN of a port	NO	YES
Configuring access VLANs	NO	YES
Enabling or disabling 802.1X authentication	NO	YES
Configuring the protocol Port-security on a port	NO	YES
Managing network devices' configuration file versioning	NO	YES
Limiting port edition rights to specific groups of users	NO	YES

Note that from the module **Dashboards**, you can gather gadgets and charts on *NetChange dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 59. Managing Network Devices

NetChange uses the SNMP protocol to query network devices and centralize all collected information in its database. You can add, import and delete network devices with an IPv4 address from the page *All network devices*. The devices can manage interfaces with IPv4 or IPv6 addresses that are displayed on the dedicated page *All addresses*. There are several ways to integrate new network devices in NetChange database:

- Adding one or several network devices using their IPv4 address.
- Importing network devices through discovery protocols (like CDP, DP or LLDP) once you added a device.
- Importing network devices using a CSV file. For more details, refer to the section [Importing Data to NetChange](#) of the chapter Importing Data.

To use NetChange at the maximum of its potential, we strongly suggest that you add at least one device using its IP address and then use the discovery protocols to add all your network devices to the page *All network devices*.

Browsing Network Devices

The network device is the highest level of organization in NetChange. All the devices that you want to manage and work with on your network are gathered on one page. You can define, check or discover how all these devices are related to each other through their VLANs, ports, routes, IP addresses and discovered items. If they support it, you can even monitor their configuration file versioning.

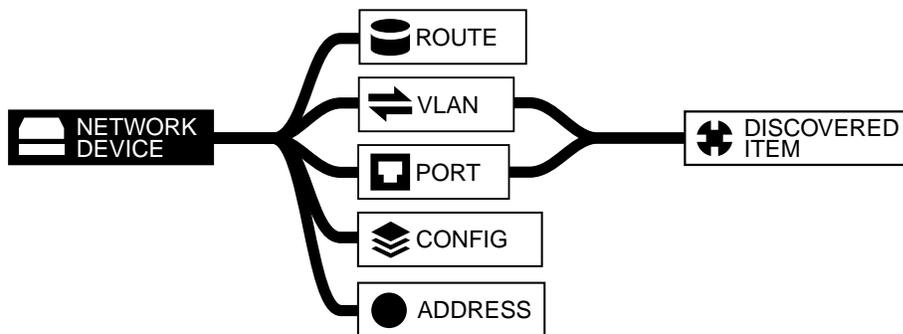


Figure 59.1. The network device in NetChange hierarchy

Browsing the Network Devices Database

To display the list of network devices

1. In the sidebar, go to **NetChange** > **Network devices**. The page **All network devices** opens.
2. You can filter the list using the column search engines.

To display a network device properties page

1. In the sidebar, go to **NetChange** > **Network devices**. The page **All network devices** opens.

- At the end of the line of the network device of your choice, click on . The properties page opens.

Browsing a Network Device Properties Page

The properties page of a network device describes its configuration and details in a set of panels.

Table 59.1. The panels on a network device properties page

Panel	Description
Main properties	The main information regarding the device. For more details, refer to the table The panel Main properties of a network device .
Refresh properties	The device refresh configuration, for both data and configuration file refresh. For more details, refer to the chapter Managing Configuration Versioning .
Additional information	All the data that is not displayed in the other panels: stack identifier, serial number, Uptime, MAC address, number of ports, etc.
SNMP properties	The supported MIBs and all the SNMP related data of the device: profile, version, port, number of retries, etc.
Configuration versioning properties	The versioning configuration status: unsupported, enabled, disabled. Once enabled and configured, the connection profile is also displayed in the panel. For more details, refer to the chapter Managing Configuration Versioning .
IP Addresses List	All the IP addresses configured on the interfaces of the device, whether IPv4 or IPv6.
Network device ports status	A graph representing the active, inactive and disabled ports of the device.
Groups access	The groups that have the network device as a resource and the rights its users have over it.

The panel **Main properties** can contain specific information, depending on the device supported MIBs, in the table below the specific properties are in italic.

Table 59.2. The panel Main properties of a network device

Property	Description
Name	The network device name.
Address	The IP address of the network device management interface.
Type	The network device type and model.
Description	The network device description and operating system.
<i>CPU Load</i>	The network device CPU Load.
Last updated	The network device last refresh date and time.
<i>Serial Number</i>	The network device serial number.
<i>Temperature</i>	The network device temperature.
Class	The network device class.

Customizing the Display on the Page All Network Devices

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Some columns provide more specific information regarding network devices:

Table 59.3. Available columns on the page All network devices

Column	Description
Complete description	All the available vendor information, the <i>SysDescr</i> of the device.
Network device OS version	The device OS version, only for Cisco equipment.
Last config. check	Device configuration file versioning information for the devices that support it.
Revision	You can also display more versioning dedicated columns. For more details, refer to the chapter Browsing the Configuration Files .
Versioning	
Analysis Time	Time necessary to collect information during the last refresh or the network device current analysis status: Being analyzed (in progress).
Uptime	The network device uptime.
Type	The network device type (level 2 switch, level 3 router...).
Network device serial number	The serial number of the network device.
Slot serial number	The slot number and slot serial number (identifier) of used slots as follows: <code><slot-number>:<slot-serial-number></code> . That column only retrieves information for used slots, empty slots are not listed.
Status	The network device status:
	<i>OK</i> : The network device is up and running.
	<i>Timeout</i> : The network device is not responding.
	<i>Misconfigured</i> : The connection with the network device is not properly set.

Adding Network Devices

The menu *Add* allows to import network devices via SNMP based on their IPv4 address, you can specify a range of IPv4 addresses to import several devices at once. The network devices can manage interfaces with an IPv4 or an IPv6 address.

Note that you can also import network devices from a CSV file. For more details, refer to the section [Importing Data to NetChange](#).

To add a network device

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. In the menu, click on **+ Add**. The wizard **Add network devices** opens.
3. If you or your administrator created classes at the network device level, in the list **Network device class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the field **IP address**, type in either the IP address of the device of your choice or the start address of a range of addresses that contains several or all of your network devices.
5. In the field **Ending IP address**, you can type in the last address of the range containing network devices.
6. You can select the SNMP profile(s) to use in order to access the SNMP agent on the devices.

Table 59.4. SNMP profile information parameters

Parameter	Description
SNMP profiles configuration	The SNMP profiles available. By default there are 3 profiles, <i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i> , but you can create as many profiles as needed to display them in this list. For more details, refer to the section Managing SNMP Profiles . Select a profile and click on  to move it to the list <i>Selected profiles</i> .
Selected profiles	This field lists the SNMP profiles to use in order to retrieve the device(s) information. SOLIDserver tries all the profiles on the device(s), following the list order. To remove a profile from the list, select it and click on  .

If you do not select any, NetChange uses the profile *standard v2c*.

7. Tick the box **Expert mode** to specify more details regarding the device(s) information retrieval. Edit the parameters according to your needs following the table below:

Table 59.5. SNMP parameters

Field	Description
SNMP port	The port used to discover the devices and retrieve their data via SNMP. By default, the port <i>161</i> is selected.
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between <i>0</i> and <i>5</i> . By default, it is set to <i>2</i> attempts.
SNMP transfer timeout (minutes)	The number of minutes above which the SNMP transfer is aborted when you add or refresh a device. You can set it between <i>0</i> and <i>999</i> . By default, it is set to <i>0</i> .
SNMP timeout	The number of seconds between each connection attempt. You can set it between <i>1s</i> and <i>5s</i> or set it to <i>10s</i> . By default, it is set to <i>5s</i> .
Use bulk	If you use SNMP version 2 or 3, you can choose to use a bulk transfer of data. This compact SNMP request method accelerates transfers by sending several requests at once. By default, it is set to <i>Yes</i> .
Use TCP	Choose to use the TCP protocol instead of the UDP when the network link is not reliable. By default <i>UDP</i> is used, the drop-down list is set to <i>No</i> .

8. In the drop-down list **Target space**, select the IPAM space associated with the network device(s).
9. Click on  to complete the operation. The report opens and works for a while before closing. The list is updated.

Once you added one device, you can retrieve all the devices it is directly connected (plugged) to thanks to the discovery protocol option detailed in the section [Importing Network Devices Using Discovery Protocols](#) below.

Importing Network Devices Using Discovery Protocols

Once you manage at least one network device, you can use the discovery protocols to retrieve and import the network devices connected to each other. For more details regarding network device addition, refer to the section [Adding Network Devices](#).

Note that you can also import network devices from a CSV file before using the discovery protocols. For more details, refer to the section [Importing Data to NetChange](#).

The discovery protocol import option retrieves all the information via three layer 2 protocols: the Cisco Discovery Protocol (CDP), the Nortel Discovery Protocol (NDP) and the Link Layer Discovery

Protocol (LLDP). The information gathered through these protocols is then retrieved using SNMP, among which, the devices neighbors i.e. the devices connected to the devices listed on the page *All network devices*.

The Cisco Discovery Protocol (CDP)

The CDP is a proprietary Data Link Layer network protocol developed by Cisco Systems to share information between devices, from their topology to their OS version, IP address or interfaces' status. NetChange uses CDP to discover Cisco network devices.

The Nortel Discovery Protocol (DP)

The DP is a Data Link Layer (OSI Layer 2) network protocol for discovery of Nortel devices and their topology. NetChange uses it to automatically discover Nortel, Avaya and Ciena network devices.

The Link Layer Discovery Protocol (LLDP)

The LLDP is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 local area network, principally wired Ethernet. LLDP is supported by the following switch vendors: HP, H3C, Nortel, Extreme Networks, Cisco and Juniper, Dell and Entreats.

To import network devices using the discovery protocols

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the device(s) for which you want to discover neighbors.
3. In the menu, select  **Import > Using CDP/NDP/LLDP**. The wizard **Add network devices** opens.
4. In the drop-down list **Target space**, select the IPAM space that should list the IP addresses of the discovered device(s).
5. Click on to complete the operation. The report opens and works for a while before closing. The devices found are listed.

The LLDP being the only vendor-neutral protocol, you might need to enable it on your devices, especially if the devices connected are from different vendors or if you connected a Nortel or Cisco device with a device from a different vendor.

Enabling LLDP on HP Devices

LLDP is enabled by default on HP Procure switches and routers. There is nothing to do. If you want to see LLDP neighbors from your HP switch, use the following command.

```
show lldp info remote-device
```

Enabling LLDP on Nortel Devices

Nortel switch 425 and 55x0 series support LLDP with a 5.x firmware. This is not enabled by default. Here is the set of command to enable LLDP:

```
5510-24T(config)#interface FastEthernet ALL
5510-24T(config-if)#lldp tx-tlv port ALL port-desc
5510-24T(config-if)#lldp tx-tlv port ALL sys-name
5510-24T(config-if)#lldp tx-tlv port ALL sys-desc
5510-24T(config-if)#lldp tx-tlv port ALL local-mgmt-addr
5510-24T(config-if)#lldp tx-tlv port ALL dot1 vlan-name ALL
5510-24T(config-if)#lldp tx-tlv port ALL dot3 link-aggregation
5510-24T(config-if)#lldp tx-tlv port ALL dot3 mac-phy
```

Depending on your firmware version, some options may be unrecognized. For VLAN, unfortunately, you need to issue the command each time you add a VLAN. When using MT, EAST or SMELT, you may want to disable ingress filtering:

```
vlan ports ALL filter-unregistered-frames disable
```

For Nortel RES 8600, there is no support for LLDP. For Nortel Switch for IBM Blade center (Nortel Layer 2-3 and 2-7), you need version 5.1 or more recent.

Enabling LLDP on Extreme Networks Devices

ExtremeOS and ExtremeWare supports LLDP with recent firmware's. You need to enable it with:

```
enable lldp ports all
configure lldp ports all advertize management-address
configure lldp ports all advertize port-description
configure lldp ports all advertize system-capabilities
configure lldp ports all advertize system-description
configure lldp ports all advertize system-name
configure lldp ports all advertize vendor-specific dot1 vlan-name
configure lldp ports all advertize vendor-specific dot3 link-aggregation
configure lldp ports all advertize vendor-specific dot3 mac-phy
```

Enabling LLDP on Cisco Devices

Starting from IS 12.2(33)SCH, LLDP is supported. Use the following command to enable it:

```
lldp run
```

On each interface, you may need to accept LLDP:

```
interface GigabitEthernet1/7
lldp enable
```

Enabling LLDP on Juniper Devices

Numerous platforms from Juniper support LLDP and LLDP-MED. The Juniper supported platforms are: EX, MX, M, J and SEX. Use the following command to enable it:

```
set protocols lldp
```

On capable and configured devices, you can see LLDP information with:

```
show lldp <detail>
```

Enabling or Disabling the 802.1X Authentication Protocol

As long as a device supports the 802.1X authentication, you can enable or disable it from the GUI. Keep in mind that:

- Enabling 802.1X authentication at device level, enables it at port level as well.

You can individually disable it on the ports. For more details, refer to the section [Managing the 802.1X Authentication Protocol on a Port](#).

- Disabling 802.1X authentication at device level, disables it at port level as well.

At port level, the protocol can still be configured but you can no longer enable it on the ports when it is disabled on their device.

We recommend that you display the column **802.1X** to see if the authentication is supported on your devices and if it is enabled or disabled. For more details, refer to the section [Customizing the List Layout](#).

To enable/disable the 802.1X authentication on a device

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Filter the list if need be.
3. Right-click on the **Name** of the device you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a network device** opens.
5. In the drop-down list **802.1x authentication**, select *Enable* or *Disable*.
6. Click on  to complete the operation. The report opens and closes. The device is marked  *Enabled* or  *Disabled* in the column **802.1X** .

Editing the SNMP Properties of a Network Device

For each device, you can set specific SNMP parameters (version, profile, port, number of retries, transfer timeout...) via the panel SNMP properties on its properties page. Keep in mind that the SNMP service and profiles must be configured beforehand. For more details, refer to the sections [Managing the SNMP Service](#) and [Managing SNMP Profiles](#).

This panel also provides an overview of the MIBs supported by the device.

To edit the SNMP properties of a network device

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the panel **SNMP properties**, click on . The wizard **SNMP parameters** opens.
4. Edit the parameters according to your needs following the table below:

Table 59.6. SNMP parameters

Field	Description
SNMP version	The version of the SNMP protocol you want to use. It can be either <i>v1</i> , <i>v2c</i> or <i>v3</i> . By default, <i>v2c</i> is selected.
SNMP port	The port that the SNMP service must use. By default, the port <i>161</i> is selected.
SNMP retries	Select the number of connection attempts when the server is in timeout. You can set it between <i>0</i> and <i>5</i> . By default, it is set to <i>2</i> attempts.
SNMP timeout	Select the number of seconds between each connection attempt. You can set it between <i>1s</i> and <i>5s</i> or set it to <i>10s</i> . By default, it is set to <i>5s</i> .
Use bulk	If you use SNMP version 2 or 3, you can choose to use a bulk transfer of data. This compact SNMP request method accelerates transfers by sending several requests at once. By default, it is set to <i>Yes</i> .
Use TCP	Choose to use the TCP protocol instead of the UDP when the network link is not reliable. By default <i>UDP</i> is used, the drop-down list is set to <i>No</i> .
SNMP transfer timeout (minutes)	Set the number of minutes above which the SNMP transfer is aborted when you add or refresh a device. You can set it between <i>0</i> and <i>999</i> . By default, it is set to <i>0</i> .

5. Click on **NEXT**. The last page of the wizard opens.
6. In the drop-down list **SNMP profile**, choose a profile using the same version of the SNMP protocol as the one you selected in the field *SNMP version*.

If you created SNMP profiles, you can choose one of your profiles. They are listed only if they use the same version of the SNMP protocol as the one you selected on the previous page.

Note that the SNMP profiles you can choose from must be configured on the appliance you are currently working with.

7. Click on to complete the operation. The wizard closes. The changes are listed in the panel.

Refreshing the Network Devices Database

After each network device import, a discovery is automatically carried out to fill NetChange database. It includes ports, VLANs, routes, IP addresses and MAC addresses information. Following this initial discovery, it is necessary to periodically refresh the database to keep it up to date. Two methods are available: a manual refresh or a scheduled refresh of network devices.

Refreshing a Device Manually

The manual refresh allows to get the latest information available regarding a network device. For instance, you should use it if its configuration or architecture have been modified.

Keep in mind that the device parameter *SNMP transfer timeout* can impact the success of the refresh. For more details, refer to the section [Editing the SNMP Properties of a Network Device](#).

To manually refresh a device

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the network device(s) you want to refresh.
3. In the menu, select  **Edit > Refresh**. The wizard **Refresh a network device** opens.
4. Tick the box **Device data** to refresh the database of the selected network device(s).
5. If you enabled versioning on the selected device(s), you can tick the box **Configuration versioning**. For more details, refer to the section [Refreshing a Configuration File](#).
6. Click on to complete the operation. The report opens and works for a while.

When the refresh is over, a report might appear and list the created IP addresses (Notice) and existing ones (Error). This list only regards the device addition or import in the selected *Target space*. You can download this report in the format of your choice: , or .

7. Click on to go back to the page **All network devices**. The page refreshes.

Scheduling a Refresh

The scheduled refresh allows to plan ahead the update of the NetChange database. You can specify different schedules depending on the devices. Typically, edge switches are queried more often than backbone routers.

Setting Up a Scheduled Refresh

The device refresh frequency can be common to several devices or specific to a device. Do not hesitate to tick one or several devices before setting up a refresh schedule.

To set up a network device scheduled refresh

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the network device(s) for which you want to schedule the refresh.
3. In the menu, select  **Edit > Scheduled refresh > Configure**. The wizard **Set refresh frequency** opens.
4. Configure the refresh frequency using the table below.

Table 59.7. Scheduled refresh parameters

Field	Description
Minute	Select the moment (o'clock, quarter past, half past or quarter to) or the frequency (in minutes) of the refresh.
Hour	Select a frequency (over the whole day or for a limited period of time each day), a set of hours or a specific hour per day for the refresh.
Date of the month	Select a specific day of the month or a frequency (every day) for the refresh.
Month	Select a specific month or a frequency (every month) for the refresh.
Day(s) of the week	Select a frequency (over the whole week or for a specific set of days) or a specific day of the week.

5. Tick the **Device data** to refresh the database of the selected network device(s) at the scheduled time.
6. If you enabled versioning on the selected device(s), you can tick the box **Configuration versioning**. For more details, refer to the section [Refreshing a Configuration File](#).
7. Click on to complete the operation. The report opens and closes. The page is visible again. The refresh frequency that you set is displayed in the panel **Refresh properties** on the properties page of your network device.

Disabling a Scheduled Refresh

Any scheduled refresh can be disabled for one or several devices at once.

To disable a network device scheduled refresh

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the network device(s) for which you want to disable the scheduled refresh.
3. In the menu, select  **Edit > Scheduled refresh > Disable**. The wizard **Disable the scheduled refresh** opens.
4. Tick the **Device data** to disable the scheduled refresh for the selected network device(s).
5. If you enabled versioning on the selected device(s), you can tick the box **Configuration versioning**. For more details, refer to the section [Refreshing a Configuration File](#).
6. Click on to complete the operation. The report opens and closes. The page is visible again.

Connecting to a Network Device Via a Console

From the properties page you can connect to a network device via an SSH, telnet or web console.

To connect to a network device via SSH

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the menu, select . **Tools** > **Connect** > **Via SSH**.
4. The SSH console connected to your device opens.

To connect to a network device via a telnet console

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the menu, select . **Tools** > **Connect** > **Via telnet**.
4. The telnet console connected to your device opens.

To connect to a network device via a web console

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the menu, select . **Tools** > **Connect** > **Via web**.
4. A new tab connecting to your device opens.

Making a Network Device Snapshot

You can retrieve information on a network device, that you manage or not, through its IP address.

EfficientIP support team might ask for a device snapshot in case of missing or distorted information on the equipment you want to add to NetChange. The snapshot is generated in .pcap format and stored in the Local files listing. Keep in mind that the SNMP service and profiles must be configured beforehand. For more details, refer to the sections [Managing the SNMP Service](#) and [Managing SNMP Profiles](#).

To make a network device snapshot

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. In the menu, select . **Tools** > **Make a snapshot**. The wizard **Make a network device snapshot** opens.
3. In the drop-down list **Interface**, select the network interface through which you want to make the snapshot.

4. In the drop-down list **SNMP profile**, select the SNMP protocol version of the snapshot generation. By default, *standard v1* is selected. For more details, refer to the section [Managing SNMP Profiles](#).
5. If you are generating a Cisco device snapshot, tick the box **Cisco device**.
6. In the field **IP address**, type in the device IP address.
7. Click on to complete the operation. The report opens and works for a while before closing. The page *All network devices* is visible again. The snapshot (*<chosen_interface>_<chosen_SNMP_profile>_snapshot.pcap* file) can be downloaded from the page **Local files listing** available from the page *Admin Home*. To download this file, refer to the procedure below.

To download a network device snapshot from the GUI

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens filtered through the bullet *Local*, under the menu.
3. In the column **Name**, the snapshot is listed and named following this format: *<chosen_interface>_<chosen_SNMP_profile>_snapshot.pcac*.
4. Filter the list if need be. Once you found the snapshot, click on its name to download it.

Creating Network Devices in Device Manager

SOLIDserver allows you to manage your network devices through NetChange and Device Manager. With a simple automated manipulation you can create, within the module Device Manager, the network devices of your choice as well as the ports and interfaces they contain.

For more details, refer to the section [Automatically Adding Network Devices in Device Manager](#).

Deleting Network Devices

If you no longer want to manage a network device and its content, you can delete it.

To delete a network device

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Tick the network device you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on to complete the operation. The report opens and closes. The device is no longer listed.

Defining a Network Device as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a network device as one of the resources of a specific group allows the users of that group to manage the network device in question as long as they have the corresponding rights granted.

Granting access to a network device as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Groups.

Chapter 60. Managing Routes

The page *All routes* is dedicated to the network devices routing tables and displays the existing routes on the layer 3 network devices you manage.

All the information displayed is retrieved via SNMP. Each route corresponds to a subnetwork and has a unique IP address and prefix. The prefix size can be any number between 0 and 32 for IPv4 addresses and between 0 and 128 for IPv6 addresses.

NetChange supports more MIBs to provide detailed route information on the page. Thanks to L3VPN you can now display the VRF routes configured on your network and also retrieve the routes Type and Protocol.

Prerequisites

- Managing network devices and/or routers that support and have SNMP enabled.
- To discover VRF routes on your network you must:
 - Make sure the routers support the MIB MPLS-L3VPN-STD-MIB.
 - Make sure the routers configured with VRF have MPLS enabled.
 - Make sure the routers configured with VRF have BGP enabled.
 - Enable the registry key `module.netchange.enable_vrf_route`.

Limitations

- Routes are retrieved automatically when you import network devices. You cannot add, edit or delete them.
- You can only export routes.
- To display the routes' *Type* and *Protocol*, the MIBs IP-FORWARD-MIB and IANA-RTPROTO-MIB must be implemented on the routers.
- Depending on your infrastructure, enabling the VRF routes discovery can significantly increase the database size.

Browsing Routes

The routes are the second level of organization in NetChange, along with the VLANs, configurations, ports and addresses.

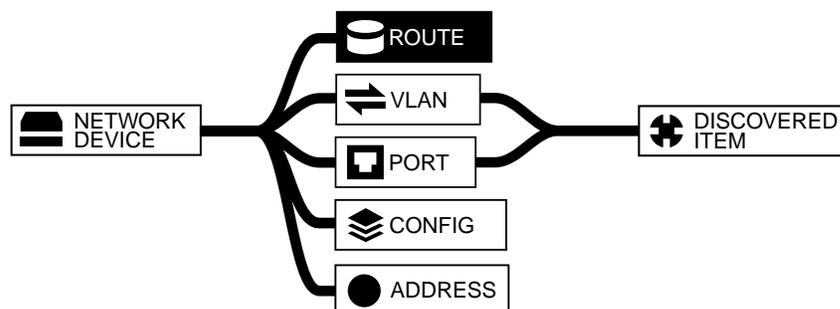


Figure 60.1. The route in NetChange hierarchy

Browsing the Routes Database

To display the list of routes

1. In the sidebar, go to  **NetChange** > **Network devices**. The page opens.
2. In the breadcrumb on the right of **Network devices**, click on  to display additional pages.
3. Click on **All routes**. The page refreshes.
4. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
5. To display the list of routes of a specific network device, in the column Network device, click on the name of the device of your choice. The page refreshes.

Understanding the Icons on the Page All Routes

Table 60.1. The icons on the page All Routes

Icon	Description on the page V4	Description on the page V6
	Routes with a prefix size between 0 to 24.	Routes with a prefix size between 0 to 96.
	Routes with a prefix size of 30.	Routes with a prefix size of 126.
	Routes with a prefix size of 31.	Routes with a prefix size of 127.
	Routes with a prefix size of 32.	Routes with a prefix size of 128.

Customizing the Display on the Page All Routes

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that you can use colored labels to differentiate at a glance IPv6 address containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

In addition to the route and network device dedicated columns, the columns **VRF name**, **VRF RD** and **Next hop** provide detailed information.

Table 60.2. Column information retrieved with L3VPN

Column	Description
VRF name	The name of the VRF as set on the network device.
VRF RD	The VRF Route Distinguisher. A unique 64-bits identifier, that can be composed of IP addresses or AS Numbers, that differentiates every set of routes on each VRF. For more details, refer to the part VRF .
Next hop	The next router to reach the prefix within the VRF.

The column **Type** provides information for each VRF.

Table 60.3. The possible values of the column Type

Value	Description
other	The routing is not working, we do not know why.
invalid/reject	The routing is not working: the route discards traffic.
local	The routing is working: it originates from a local interface.

Value	Description
remote	The routing is working: it has a remote destination.
blackhole	The routing is working: the route is discarding traffic silently.

The column **Protocol** allows to determine how the VRF was found.

Table 60.4. The possible values of the column Protocol

Value	Description
other	No protocol is specified.
local	Local interface.
netmgmt	Static route.
icmp	Result of ICMP Redirect.
egp	Exterior Gateway Protocol.
ggp	Gateway-Gateway Protocol.
hello	FuzzBall HelloSpeak
rip	Berkeley RIP or RIP-II.
is-is	Dual IS-IS.
es-is	ISO 9542.
ciscoigrp	Cisco IGRP.o
bbnSpfIgp	BBN SPF IGP.
ospf	Open Shortest Path First.
bgp	Border Gateway Protocol.
idpr	InterDomain Policy Routing.
ciscoEigrp	Cisco Enhanced Interior Gateway Routing Protocol.
dvmrp	Distance Vector Multicast Routing Protocol.
rpl	Routing Protocol for LLNs [RFC-ietf-roll-rpl-19].
dhcp	Dynamic Host Configuration Protocol [RFC2132].
ttdp	Train Topology Discovery Protocol [IEC 61375-2-5].

Enabling the Registry Key Required to Display the VRF Routes

Once you met the [Prerequisites](#), you can enable a registry key to retrieve all the routes of all VRFs configured on your network devices providing L3VPN services.

Enabling the key updates the data on the page *All routes*, including additional information in the columns *VRF name* and *VRF RD*.

To enable the registry key that controls the switch based on time drift

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *module.netchange.enable_vrf_route*. The entry is the only one listed.

4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in *1* to enable it. By default, its value is *0*.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Chapter 61. Managing VLANs

The page *All VLANs* provides an overview of the existing Virtual Local Area Networks for each network device and their ID if you purchased the license NetChange-IPL. If you have the NetChange license, it also allows to add, edit and delete VLANs on your devices. For more details regarding the two available NetChange licenses, refer to the table [NetChange licenses differences](#).

Browsing VLANs

The VLANs are the second level of organization in NetChange, along with the routes, configurations, ports and addresses.

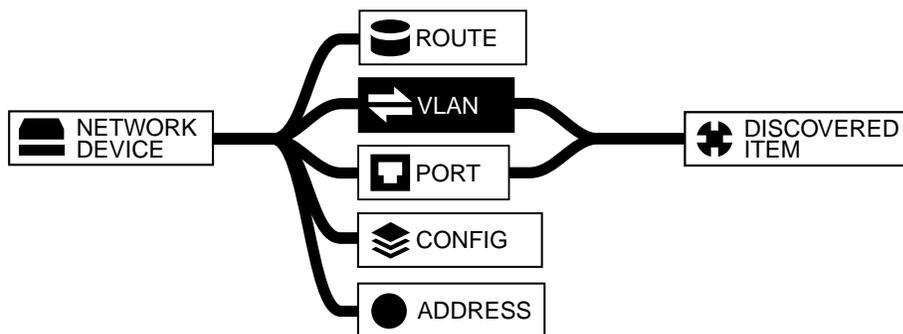


Figure 61.1. The VLAN in NetChange hierarchy

Browsing the VLANs Database

To display the list of VLANs

1. In the sidebar, go to **NetChange > VLANs**. The page **All VLANs** opens.
2. To display the list of VLANs of a specific network device, in the column **Network device**, click on the name of the device of your choice. The page refreshes.

To display a VLAN properties page

1. In the sidebar, go to **NetChange > VLANs**. The page **All VLANs** opens.
2. At the end of the line of the VLAN of your choice, click on **ⓘ**. The properties page opens.

Customizing the Display on the Page All VLANs

Users of the group *admin* can create customized column layouts. The button **☰ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that the column **Port list** contains the number of all the ports associated with each VLAN. You can edit this list if you purchased the license NetChange, otherwise this list is merely informative.

Adding a VLAN

With the NetChange license, you can add VLANs to the page *All VLANs* and then associate them with existing ports. Using 802.1q VLAN Trunking protocol, a VLAN can cover a network area on multiple switches.

To add a VLAN from the page *All VLANs*

1. In the sidebar, go to  **NetChange** > **VLANs**. The page **All VLANs** opens.
2. In the menu, click on **+ Add**. The wizard **Add a VLAN** opens.
3. In the field **Name**, name the VLAN.
4. In the field **VLAN ID**, type in an ID between 1 and 1005 for your VLAN.
5. In the drop-down list **Network device**, select the network device where you want to add your VLAN.
6. Click on to complete the operation. The report opens and closes. The VLAN is listed.

You can also add a VLAN from the list *All VLANs* of a specific device, in this case the *Network device* drop-down list does not appear.

In addition, you can use existing VLANs ID and name and add them to another device. That way, you only need to specify a device for the VLAN name and ID to be used automatically upon creation. Obviously, the ports configuration of the selected VLAN is not created in the target network device.

To add a VLAN from the page *All VLANs* using an existing name and ID

1. In the sidebar, go to  **NetChange** > **VLANs**. The page **All VLANs** opens.
2. Right-click over the **ID** of the VLAN of your choice. The contextual menu opens.
3. Click on . The **Add a VLAN** wizard opens.
4. In the fields **Name** and **VLAN ID** are displayed in gray the name and ID of the chosen VLAN.
5. In the drop-down list **Network device**, select the device of your choice.
6. Click on to complete the operation. The report opens and closes. There are now two VLAN with the same name and ID listed, only their device differs.

Editing a VLAN

Editing a NetChange VLAN means renaming it. However, with the NetChange license you can decide to use it with one or several of your network ports. For more details regarding the port and VLAN interaction, refer to the section [Associating a Port With a VLAN](#).

To rename a VLAN

1. In the sidebar, go to  **NetChange** > **VLANs**. The page **All VLANs** opens.
2. Filter the list of need be.
3. Right-click over the **ID** of the VLAN you want to rename. The contextual menu opens.
4. Click on . The **Add a VLAN** wizard opens.
5. In the field **Name**, rename the VLAN.

6. In the field **VLAN ID**, the ID is displayed but cannot be edited.
7. Click on to complete the operation. The report opens and closes. The list refreshes, the new VLAN name is listed.

Deleting a VLAN

With the NetChange license you can delete any VLAN from any network device as long as it is not used on any port.

To delete a VLAN

1. In the sidebar, go to  **NetChange > VLANs**. The page **All VLANs** opens.
2. In the column **Network device**, click on the name of the device of your choice to display only its VLANs.
3. Tick the VLAN(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The VLAN is no longer listed.

Chapter 62. Managing Ports

The ports are physical interfaces of the network devices. NetChange discovers the network devices ports using a discovery algorithm that automatically analyzes each port and displays its type and status. It also allows to know which MAC or IP addresses should be looked for and the devices connection on the network. Typically the listed ports can be:

- **Edge or terminal ports:** used to connect the terminal network devices of the network (servers, workstations, printers, ...);
- **Interconnection ports:** used to link the network devices between them (the backbone).

Depending on your network devices, some ports can actually be both. Some columns on the page provide all this information:

- **Interco** (for interconnection) is purely informative even if you can manually force its value to *Yes*, *No* or *Autodetect* in the GUI.
- **Trunking/Tagging mode** provides the actual port type, edge ports are marked *Access* and interconnection ports are marked *Trunk* or *Tagged*.

NetChange module allows to edit a port and associate it with existing VLANs on your device (existing by default or that you added). To be able to edit a port, you must meet the following **prerequisites**:

1. The SNMP service is configured properly. For more details, refer to the section [Managing the SNMP Service](#).
2. The SNMP profile used with network devices has a read/write access to interact with the device(s). For more details, refer to the sections [Managing SNMP Profiles](#) and [Editing the SNMP Properties of a Network Device](#).
3. You have the NetChange license. NetChange-IPL does not provide port edition options. For more details, refer to the table [NetChange licenses differences](#).
4. The network device on which you edit the port supports MIBs that allow port edition.

Once these prerequisites are met, you can edit your ports. This allows to associate them with any VLAN on your network or even use them in a tagged or untagged mode and influence their behavior on the network. As a general rule, when choosing to tag or not a port you should take into account the following:

- The untagged mode (called *Access* on Cisco devices) uses the ID of the tagged VLAN the port is associated with when sending and receiving data. That way packages are identified throughout the transfer on the network from the sending port to the receiving one. Once the package is received, the tag number is dismissed, in other words, untagged. This transfer mode is based on terminal, or edge, ports as packages always reach their destination thanks to their tag once sent. In the columns *VLAN name list* and *VLAN # list*, the untagged/access VLAN of the port is followed by a star.
- The tagged mode (called *Trunk* on Cisco devices) uses the ID of the VLANs associated with the port only when sending packages. The tag identifies the target port. Once the package is received, the tag number is kept. This transfer mode is based on interconnection ports as it allows to send out data all over the network.

Browsing Ports

The ports are the second level of organization in NetChange, along with the routes, configurations, VLANs and addresses.

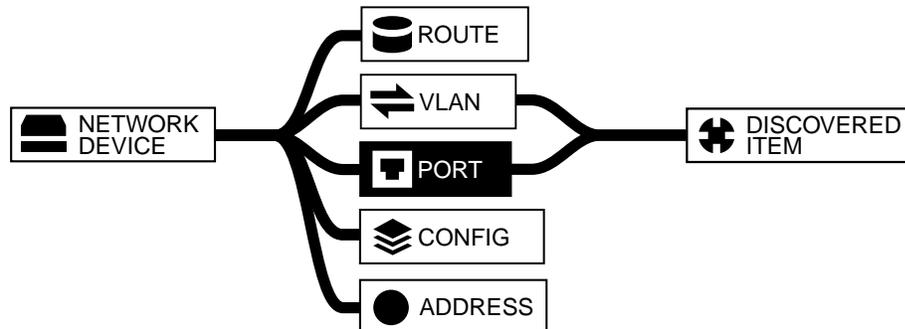


Figure 62.1. The port in NetChange hierarchy

Browsing the Ports Database

To display the list of ports

1. In the sidebar, go to **NetChange > Ports**. The page **All ports** opens.
2. To display the list of ports of a specific network device, in the column Network device, click on the name of the network device of your choice. The page refreshes.

To display a port properties page

1. In the sidebar, go to **NetChange > Ports**. The page **All ports** opens.
2. At the end of the line of the port of your choice, click on . The network port properties page opens.

Customizing the Display on the Page All Ports

Users of the group *admin* can create customized column layouts. The button **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that the columns can help you filter the list on specific port settings: *Trunking/Tagging mode*, *Configured speed*, *Configured duplex*, *VLAN name list*, *VLAN # list*, *Operating speed*, *Operating duplex...*

Understanding the Port Statuses

The column **Status** provides information regarding the ports you manage.

Table 62.1. NetChange port statuses

Status	Description
✔ <i>Active</i>	The port is active, or up.
✘ <i>Inactive</i>	The port is inactive, or down.
✔ <i>Testing</i>	The port is up but no operational packets can be passed.
✘ <i>LowerLayerDown</i>	The port is inactive. These statuses are very rare. For more details, refer to the description of the MIB <i>IF-MIB</i> .
✘ <i>NotPresent</i>	
✘ <i>Dormant</i>	
❗ <i>Unknown</i>	The port status is unknown.
❗ <i>Disabled</i>	The port was disabled. For more details, refer to the section Enabling or Disabling a Port .

Enabling or Disabling a Port

As an experiment, NetChange offers a function to disable or enable ports of any network device, as soon as it has been validated. This function allows you to disable directly through the web interface any port, typically when a workstation has been detected as infected by a virus or when a user has not been authorized on the network. To work properly, you must have defined the *Write community* string of the SNMP profile used by the network device.

Keep in mind that **you should never disable interconnection ports** as you take the risk of losing access to your network device.

To enable/disable a port

1. In the sidebar, go to  **NetChange** > **Ports**. The page **All ports** opens.
2. Tick the port(s) for which you want to change the status.
3. In the menu, select  **Edit** > **Port status** > **Enable** or **Disable**. The wizard **Change the status of a port** opens.
4. Click on  to complete the operation. The report opens and closes. The page is visible again.

Keep in mind that on some devices, especially Cisco Catalyst, the configuration is not written after the modification has been done, so if no *write configuration* command is made through CLI, modifications can be lost if the switch is reloaded.

Editing a Port Interconnection

Interconnection ports are used to link the network devices (the backbone) between them. Most of the network traffic is done through these ports.

NetChange discovery algorithm automatically isolates interconnection ports: they are marked *Yes* in the column *Interco* if the number of discovered items on a port is greater than a defined limit (4 by default).

The value of the column *Interco* is merely a way of filtering the ports in the list. However, if a port interconnection is set to *Yes*, any MAC address discovered via the port is not logged on the page *Discovered item history*. For more details, refer to the section [Tracking the Discovered Items History of a Specific Device](#).

To tag a port interconnection status

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Tick the port(s) you want to use as an interconnection.
3. In the menu, select  **Edit > Interconnection**. The wizard **Editing the port interconnection status** opens.
4. In the drop-down list **Port interconnection status**, select:
 - a. **Auto** to automatically detect the interconnection status of the port.
 - b. **Yes** to force the port interconnection status to *Yes*.
 - c. **No** to force the port interconnection status to *No*.
5. Click on to complete the operation. The report opens and closes. The page is visible again.

If you have selected *Auto*, the value of the **Interco** column is **Yes** or **No** depending on the interconnection status of the port.

If you have selected *Yes* or *No*, the value of the **Interco** column switches to **Yes (forced)** or **No (forced)** respectively.

Editing a Port Speed and Duplex Mode

You can edit the port speed and duplex on each port individually. In some cases you might only be able to edit the speed. During the configuration, the available values depend on the port possible speed and duplex configuration. We recommend that you display the **Configured speed**, **Configured duplex**, **Operating speed** and **Operating duplex** columns to rapidly see the speed and duplex configuration of the ports.

Keep in mind that you can only see the speed and duplex changes of active ports. If you edit the port and speed of an inactive port, the changes are never visible in the GUI.

To edit a port speed and duplex mode

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Filter the list if need be.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on . The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the drop-down list **Speed and duplex mode**, select the port *<speed> <duplex>* of your choice.
7. Click on **NEXT**. The last page of the wizard opens.
8. Click on **OK** to complete the operation. The report opens and closes. The ports list is visible again.

Once you edited the port speed and duplex, you need to refresh the port to see your changes in the columns **Configured speed** and **Configured duplex**. For more details, refer to the section [Refreshing the Ports Database](#) below.

Updating a Port Description

Even though the name of a switch port cannot be edited, it is possible to modify its description through NetChange to help you recognize it instantly from the graphical interface. The description is directly updated on the port itself and is visible by any user that discovers the device. To work properly, you must have defined the *Write community* of the SNMP profile used by the network device.

To update a port description

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Tick the port(s) for which you want to change the description.
3. In the menu, select  **Edit > Update port description**. The wizard **Update port description** opens.
4. In the field **Port description**, type in the description of your choice.
5. Tick the box **Refresh NetChange** if you want to refresh the ports list immediately after your modification.
6. Click on **OK** to complete the operation. The report opens and closes. The page is visible again and displays the new value in the column **Description**.

Managing the 802.1X Authentication Protocol on a Port

You can manage 802.1X authentication on a port as long as you meet the [prerequisites](#) and take into account the [limitations](#).

Prerequisites

- The license *NetChange*.
- The network device that the port belongs to must support 802.1X authentication.
- 802.1X authentication must be enabled on the device to be managed from the GUI.
- 802.1X authentication must be configured on each port individually.
- 802.1X authentication must be enabled at device level to be enabled at port level. If it is disabled on the device, the port is configured with it but not used for authentication.

Limitations

- On HP devices, only the HP-DOT1X-EXTENSIONS-MIB is supported.

- On Cisco devices, the interface *vlanTrunkPortDynamicState* should not be set to *auto* or *desirable*.

Enabling or Disabling 802.1x Authentication Protocol on a Port

You must edit ports to enable or disable 802.1X authentication. You can disable 802.1X authentication on ports even if it is enabled on the network device they belong to.

We recommend that you display the column **802.1X** on the pages *All network devices* and *All ports*. For more details, refer to the section [Customizing the List Layout](#).

To enable/disable the 802.1X authentication on a port

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Click on the name of a device marked *Active* in the column **802.1X**. The page **All ports** of the device opens.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The last page of the wizard opens.
7. The section **Security configuration** contains the box **802.1X**. It is grayed out if 802.1X authentication is disabled on the device.
 - a. To enable the 802.1X authentication, tick the box. The page refreshes.
 - b. To disable the 802.1X authentication, untick the box. The page refreshes.
8. Click on **OK** to complete the operation. The report opens and closes. The port is marked  *Enabled* or  *Disabled* in the column **802.1X**.

Note that you can enable both 802.1x authentication and port-security on a port, but on HP devices, if 802.1x is enabled, only the port-security mode *802 1xAuthorised* is available.

Restricting Access to a Port with the Protocol Port-security

You must edit ports to use the protocol Port-security if the device supports it.

This protocol allows to restrict input to an interface by limiting and identifying MAC addresses that are allowed to access the port.

By default, the protocol is enabled on the devices that support it and you can enable or disable it individually on each port.

Prerequisites

- The license *NetChange*.

- On Cisco devices, the port mode Trunking/Tagging (i.e. *switchport*) is set to *Access* or *Trunk*.

Limitations

- On Cisco devices, only the CISCO-PORT-SECURITY-MIB is supported.
- Only HP devices supporting the HP-ICF-GENERIC-RPTR MIB can be configured.
- On HP devices, if you enable Port-security and 802.1x on a port, only the port-security mode *8021xAuthorised* can be configured.
- On HP devices, if you enable Port-security and 802.1x on a port, you cannot limit the number of MAC addresses that can access the port.

Enabling or Disabling Port-Security

To enable or disable Port-security on a port, you must edit it. You can enable both port-security and 802.1x authentication on a port, but keep in mind that on HP devices, if 802.1x is enabled, only the port-security mode *8021xAuthorised* is available.

We recommend that you display the columns **Port-security** on the page *All ports* to see which ports support it. For more details, refer to the section [Customizing the List Layout](#).

To enable/disable the Port-security option on a port

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Click on the name of a device of your choice. The page **All ports** of the device opens.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The last page of the wizard opens.
7. In the section **Security configuration**:
 - a. To enable Port-security, tick the box **Port-security**. The page refreshes and displays the fields *Mode*, *Maximum number of secured MAC addresses* and/or *Action*. For more details, refer to the section [Configuring Port-Security Modes on HP or Cisco Devices](#) and or [Limiting the Number of MAC Addresses Allowed to Access a Port](#).
 - b. To disable Port-security, untick the box **Port-security**. The page refreshes, the dedicated fields are no longer visible.
8. Click on **OK** to complete the operation. The report opens and closes. The port is marked  *Enabled* or  *Disabled* in the column **Port-security**.

Limiting the Number of MAC Addresses Allowed to Access a Port

You can limit access to a port by setting a maximum number of MAC addresses when you edit the port. This option can only be set if you enable the protocol Port-security.

We recommend that you display the columns **MAC number limit** on the page *All ports* to see which ports you restricted the access to, by default it contains the value *1*. For more details, refer to the section [Customizing the List Layout](#).

Note that on HP devices you cannot set a maximum number of MAC address on a port configured both with 802.1x authentication and Port-security, the Port-security mode *8021xAuthorised* does not allow it.

To set a maximum number of MAC addresses that can access a port

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Click on the name of a device of your choice. The page **All ports** of the device opens.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The last page of the wizard opens.
7. In the section **Security configuration**, tick the box **Port-security**. The page refreshes.
8. If you are editing the port of a HP device, make sure the Port-security **Mode** selected is either *FirstN* or *LimitedContinuous*.
9. In the field **Maximum number of secured MAC addresses**, type in the number of MAC addresses that can access the port. This number depends on your device. By default, Port-security is configured with *1* MAC address.

To configure an *Action* for the port once the number of secured MAC addresses exceeds the one you just set, refer to the section [Configuring Port-Security Modes on HP or Cisco Devices](#).

10. Click on **OK** to complete the operation. The report opens and closes. The column **MAC number limit** displays the new value.

Configuring Port-Security Modes on HP or Cisco Devices

SOLIDserver supports the configuration of some Port-security modes on the ports belonging to HP and Cisco devices.

To configure the option Port-security on HP devices

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Click on the name of the HP device of your choice. The page **All ports** of the device opens.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The last page of the wizard opens.
7. In the section **Security configuration**, tick the box **Port-security** if it is not ticked.
8. In the field **Mode**, select *FirstN*, *8021xAuthorized* or *LimitedContinuous*. If the authentication 802.1x is enabled, you can only select *8021xAuthorized*.

Note that the modes *FirstNConditionally* and *configureSpecific* cannot be selected during the Mode configuration. They are automatically retrieved and displayed if they are configured on a port you are editing.

9. In the field **Action**, select *disable*, *sendTrap* or *sendTrapAndDisablePort*.

If you configured a *Maximum number of secured MAC addresses*, the *Action* applies to the extra MAC addresses: it is performed when the number of MAC addresses exceeds the one you set in the field.

10. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

To configure the option Port-security on Cisco devices

1. In the sidebar, go to **NetChange > Network devices**. The page **All network devices** opens.
2. Click on the name of a device of your choice. The page **All ports** of the device opens.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on **✎**. The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Click on **NEXT**. The last page of the wizard opens.
7. In the section **Security configuration**, tick the box **Port-security** if it is not ticked.
8. In the field **Action**, select *shutdown*, *dropNotify* or *drop*.

If you configured a *Maximum number of secured MAC addresses*, the *Action* applies to the extra MAC addresses: it is performed when the number of MAC addresses exceeds the one you set in the field.

9. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Limiting Port Edition Rights to Specific Groups of Users

With the license NetChange, you can use a class and a rule to define which groups of users are allowed to edit the ports of your choice. For instance, you can decide that all the interconnection ports can only be edited by users of the group *admin*.

To limit edition rights on a port you must:

1. Enable the class *reserved* in Class Studio.
2. Add and configure the rule 378.
3. Apply the class to the port you want to limit access to.

Keep in mind that once you configured the rule and enabled and applied the class to a port, only the users belonging to a group of users authorized in the rule can edit that port. Even users with edition permissions on the ports and the proper resources are not allowed to edit the ports configured with the class *reserved* if they do not belong to an authorized group of users.

Enabling the Class reserved

By default, the class *reserved* is disabled. You must enable it to be able to apply it to a port and restrict access to it once you configured the rule 378.

To enable the class reserved

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the column **Type**, type in *Port* and hit Enter.
4. In the column **Name**, type in *reserved* and hit Enter.
5. Tick the class *reserved* that does not belong to the directory Library.
6. In the menu, select  **Edit > Enable class** or **Disable class**. The wizard opens.
7. Click on to complete the operation. The report opens and closes, the page refreshes. The class is marked as  *Enabled* in the column **Status**.

Adding and Configuring the Rule 378

You must add and configure the rule 378 to be able to use the class *reserved* to restrict access to your ports.

Keep in mind that once the rule is configured, its configuration overrides the port permissions set for any group of users. Even if a group has all the existing ports within its resources and edition permissions over them, its users are not allowed to edit the ports set with the class *reserved* if they are not listed among the *Authorized groups*.

To add the rule 378 that defines the groups of users allowed to edit reserved ports

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the menu, click on  **Add**. The **Add a rule** wizard opens.
4. In the drop-down list **Module**, select *NetChange*.
5. In the drop-down list **Event**, select *Edit: ports properties*.
6. In the list **Rule**, select *(378) Ports configuration permission*.
7. In the field **Rule name**, name the rule. That name is listed in the column *Instance*.
8. In the field **Comment**, you can type in a comment if you want.
9. Click on . The page **Rule filters** opens.
10. Click on . The page **Rule parameters** opens.

11. In the list **Available groups**, select one by one the group(s) of users that can edit the ports configured with the class *reserved*.
12. Click on . The group is moved to the list **Authorized groups**. Repeat this action for as many groups as needed.
13. Click on  to complete the operation. The report opens and closes. The rule is listed.

Applying the Class reserved to a Port

Once you enabled the class *reserved* and configured the rule 378, you can apply the class to the ports of your choice.

We recommend that you display the column **Class** on the page All ports to have an overview of the ports you applied it to. For more details, refer to the section [Customizing the List Layout](#).

To apply the class reserved to a port

1. In the sidebar, go to  **NetChange** > **Ports**. The page **All ports** opens.
2. Filter the list if need be.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. In the list **Port class**, select *reserved*.
6. Click on  until you get to the last page of the wizard.
7. Click on  to complete the operation. The report opens and closes. The ports list is visible again.

Once the class is applied on a port, only users belonging the *Authorized groups* configured in the rule can edit the port.

Configuring VLAN Tagging on a Port

With the license *NetChange*, you can edit the ports configuration on the network according to your needs. For instance, set a port with a specific access VLAN.

Keep in mind that the untagged/access VLAN tag of a port is followed by a star (*) in the columns *VLAN name list* and *VLAN # list*.

Configuring the Tagging Mode

When configuring VLAN tagging at ports level, you must choose the relevant tagging mode before associating a port of specific VLANs. For this reason, we recommend that you display the **Trunking/Tagging mode** column to rapidly see your ports configuration.

There are different tagging modes available depending on the network device vendor: Cisco or others.

Cisco devices tagging modes

Cisco devices offer three tagging modes: Trunk (i.e. tagged), Access (i.e. not tagged) and auto (the port mode is automatically one or the other). Setting a port to Trunk mode sets its tag encapsulation mode to 802.1Q .

Other vendors tagging modes

Non-Cisco devices offer two tagging modes: tagged or mixed.

Keep in mind that you can only edit a device Trunking/Tagging mode if the SNMP configuration set at device level allows to retrieve the MIBs.

To edit a port tagging mode on a non-Cisco network device

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Filter the list if need be.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the drop-down list **Trunking/Tagging mode**, select the mode of your choice: *Tagged* or *Mixed*.
7. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Once you edited the port tagging mode, you need to refresh the port to see your changes in the column **Trunking/Tagging mode**. For more details, refer to the section [Refreshing the Ports Database](#) below.

To edit a port tagging mode on a Cisco network device

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Filter the list if need be.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the drop-down list **Trunking/Tagging mode**, select the mode of your choice: *Trunk*, *Access* or *Auto*.

If you set the trunking/tagging mode to *auto*, the 802.1X authentication must be inactive.

7. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Once you edited the port tagging mode, you need to refresh the port to see your changes in the column **Trunking/Tagging mode**. For more details, refer to the section [Refreshing the Ports Database](#) below.

Associating a Port With a VLAN

Through the port edition wizard you can associate a port with a set of VLANs. These VLANs can be tagged or untagged depending on the port tagging mode.

To rapidly see your port/VLAN association, make sure the columns **VLAN # list** and **VLAN name list** are displayed.

Associating a Port with an Untagged VLAN

To associate a port with an untagged VLAN, its mode must be Access or Auto (on Cisco devices) or Mixed (on any other device vendor). To edit the port tagging mode, refer to the section [Configuring the Tagging Mode](#) above.

You can only associate one untagged VLAN with a port.

To associate a port with an untagged VLAN

1. In the sidebar, go to  **NetChange** > **Ports**. The page **All ports** opens.
2. Filter the list if need be.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the section **VLAN addition**, the drop-down list **Access/Untagged VLAN** displays the untagged VLAN associated with your port. Select the VLAN of your choice. By default, the *1 - default* VLAN is selected.

If your port mode is *Mixed* or *Auto*, the previously selected VLAN is moved to the list **Available VLANs**.

The lists can be empty if the device port cannot be configured.

7. Click on **OK** to complete the operation. The report opens and closes. The untagged VLAN associated with the port is followed by a * in the columns **VLAN # list** and **VLAN name list**.

Associating a Port with a Tagged VLAN

To associate a port with tagged VLANs, its mode must be Trunk or Auto (on Cisco devices) or Tagged or Mixed (on any other device vendor). To edit the port tagging mode, refer to the section [Configuring the Tagging Mode](#) above.

You can add as many untagged as you want with a port.

To associate a port with a tagged VLAN

1. In the sidebar, go to  **NetChange** > **Ports**. The page **All ports** opens.
2. Filter the list if need be.

3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port** opens.
5. If you or your administrator created classes at the port level, in the list **Port class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the section **VLAN addition**, the field **Trunk/Tagged VLAN list** displays all the tagged VLANs associated with your port.

To add tagged VLANs, in the list **Available VLANs**, select a VLAN and click on , or double-click on it, to move it to the list **Trunk/Tagged VLAN list**.

To remove the port association with tagged VLANs, in the list **Trunk/Tagged VLAN list**, select a VLAN and click on , or double-click on it, to move it back to the list **Available VLANs**.

7. Click on **OK** to complete the operation. The report opens and closes. The VLANs associated with the port are displayed in the columns **VLAN # list** and **VLAN name list**.

Refreshing the Ports Database

SOLIDserver allows the scheduled refresh of network devices - for more details, refer to the section [Scheduling a Refresh](#) -. Still, you have the possibility to manually refresh the information for a selection of ports of a specific device directly from the page All ports.

To refresh the ports information manually

1. In the sidebar, go to  **NetChange > Ports**. The page **All ports** opens.
2. Filter the list using the column **Network device** if need be.
3. Tick the port(s) you want to refresh.
4. In the menu, select  **Edit > Refresh**. The wizard **Refresh a network device** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Chapter 63. Managing Configuration Versioning

You can manage the configuration file versioning of network devices. Versioning allows to automatically save all the changes in the configuration files, and all the revisions of the files are saved in SOLIDserver backup file. To monitor versioning from NetChange you must:

1. Create a connection profile from the Administration module.
2. Enable versioning and configure the connection profile on the device.

Once configured, all the revisions are listed on the page **All configurations**, and you can compare changes on the page **Configuration** and download entire files or only changes.

Prerequisites

- SOLIDserver with the license NetChange.
- The network device must support versioning, like HP, Cisco, Juniper or Extreme.
- The network device must already be managed from SOLIDserver. To add a network device, refer to the section [Adding Network Devices](#).

Limitations

- Disabling the configuration files versioning on a device deletes all the configuration file revisions from the database.
- Configurations files are local. You cannot push them to the network device you manage from SOLIDserver.
- Configurations files are read-only. You cannot rollback to a former file version.
- On appliances configured with High Availability, the configuration files versioning of your network devices can only be managed from the Master appliance.

Browsing Configuration Files

The configurations are the second level of organization in NetChange, along with the routes, VLANs, ports and addresses. They outline the devices use and detail their structure.

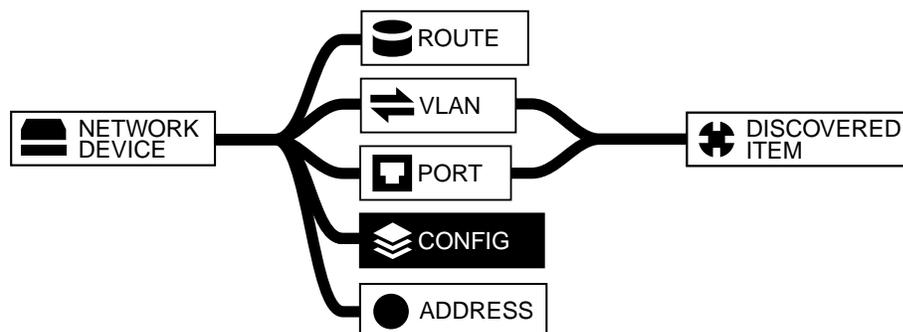


Figure 63.1. The configuration in NetChange hierarchy

Browsing the Configuration Files Database

The page **All configurations** allows to browse all the revisions of the configuration file of the network devices for which you enabled versioning. The page **Configuration** displays the content of the configuration file revision of your choice.

To display the list of configuration files

1. In the sidebar, go to  **NetChange > Configurations**. The page **All configurations** opens.
2. To display the list of configuration files of a specific device, in the column **Name**, click on the name of the network device of your choice. All the revisions of the configuration files of the selected device are displayed.

There are 6 columns on the page **All configurations** that you can sort and filter. By default, **all the columns are displayed on the page**, you cannot change their order.

Table 63.1. The columns on the page All configurations

Column	Description
Name	The network device name. This column is not displayed if you list the configurations files of a specific network device.
Vendor	The network device vendor. This column is not displayed if you list the configurations files of a specific network device.
IP address	The IP address of the network device management interface.
Revision	The configuration file revision number, the latest revision is displayed as follows: r.<number> (last) . This revision number is based on a version control system and does not reflect the number of changes in the file.
Configuration date	The date and time when the configuration was saved. It is unique to each revision.
Changed lines	The number of lines that changed between the latest revision and the previous one.

At the end of the line of each revision listed, you can find two links.

Table 63.2. The links on the page All configurations

Link	Description
Compare with previous	This link allows to compare side by side the latest revision of the configuration file with the previous one.
Download	This link allows to download any revision of the configuration file of your choice. The downloaded file is named <code><network-device-ip-address>_r<revision-number>.txt</code> .

Do not hesitate to set up alerts on the page to monitor any changes. For more details, refer to the chapter [Managing Alerts](#).

For more details regarding downloads, refer to the section [Downloading Versioning Information](#).

To display the latest configuration file from the page All network devices

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. In the column **Revision**, click on **r.<revision-number>**. The page **Configuration** opens and displays the latest version of the network device configuration file.

Above the configuration file content, a gray area contains the network device details: its name, its IP address, and the date and time when the configuration version was saved.

To display a configuration file from the page All configurations

1. In the sidebar, go to  **NetChange > Configurations**. The page **All configurations** opens.
2. Filter the list if need be. If you click on a network device **Name**, only the revisions of that device are displayed.
3. In the column **Revision**, the latest revision number is named *r.<number> (last)*, click on the revision number of your choice. The page **Configuration** opens and displays the chosen revision of the configuration file.

Above the configuration file content, a gray area contains the network device details: its name, its IP address, and the date and time when the configuration version was saved.

By default, the configuration file displays all the passwords in plain text, if you do not want all the users to see them you can edit a registry database entry to hide them. For more details, refer to the section [Configuring the Passwords Display in the Configuration Files](#).

Displaying Versioning Information on the page All network devices

A set of columns dedicated to network devices versioning provide information directly from the page **All network devices**.

Table 63.3. The columns on the page All network devices

Column	Description
Last config. check	The date and time of the last configuration check on the device. This column is displayed by default on the page.
Revision	The last known device configuration revision number, following the format r.<revision-number> . This revision number does not reflect the number of changes on the device configuration file. The number is a link toward the latest version of the configuration file on the page <i>Configuration</i> . This column is displayed by default on the page.
Versioning	The device configuration versioning status: Yes means the versioning is active; No means it is inactive; Error means it is active but an error occurred during the last operation; Unsupported means that versioning is not supported by the device. This column is displayed by default on the page.
Connection profile	The device connection profile, the name of the profile currently associated with the device.
Configuration refresh	The configuration refresh schedule for the versioning.
Last configuration	The date and time of the last configuration date, the last known configuration for the device. It corresponds to the <i>Revision</i> number displayed on the page <i>All network devices</i> .
First configuration	The date and time of the first configuration date, the first known configuration for the device. It corresponds to the first known <i>Revision</i> available on the page <i>All configuration</i> for the device.
No config. change for	The number of hours that passed between the <i>Last config. check</i> and the <i>Last configuration</i> on the device.

Do not hesitate to set up alerts based on these dedicated columns. For more details, refer to the chapter [Managing Alerts](#).

Managing Connection Profiles

To enable and configure versioning on a network device, a connection profile must be created and associated with a network device that supports it.

Connection profiles are managed on the page **Network devices connection profiles** in the module Administration.

Only from the group *admin* can add, edit and delete connection profiles.

Adding Connection Profiles

They can create as many connection profiles as needed.

Keep in mind that one profile can be used for several devices as long as the profile configuration suits all the devices it is associated with.

To add a connection profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.
3. In the panel **Network devices connection profiles**, click on **ADD**. The wizard **Add a connection profile** opens.
4. In the field **Connection profile name**, name the profile.
5. In the field **Description**, you can type in a description for the profile.
6. In the drop-down list **Connection protocol**, select the protocol used by the profile: *SSH*, *TELNET* or *RSH*.
7. In the field **Login**, you can type in the name of the user connecting.
8. In the field **Password**, you can type in the connecting user password.
9. If you want to configure the profile with more details, tick the box **Expert mode** and follow the table below.

Table 63.4. Connection profile Expert mode fields

Field	Description
"enable" login	You can type in the name of user with enable privileges. This field can empty. In connection profiles created for Juniper network devices, you cannot use a superuser login.
"enable" password	You can type in a password of the enable mode on the network device. This field can be empty. In connection profiles created for Juniper network devices, you cannot use a superuser login.
Auto enable	You can tick this box if the enable mode is automatic on the network device.
No "enable" privileges	You can tick this box if you do not want to use enable privileges for the profile. In connection profiles created for Juniper network devices, you might have to tick this box.
Cypher type	If you selected <i>SSH</i> as your connection protocol, you can select the protocol cypher type: <i>3des</i> , <i>blowfish</i> or <i>des</i> . By default, <i>3des</i> is selected.
SSH identity	If you selected <i>SSH</i> as your connection protocol, you can paste your SSH key.
Port	If you selected <i>TELNET</i> as your connection protocol, you can specify the Telnet port number on the network device.
Connection timeout	You can set the number of seconds after which the network device is considered unreachable. By default, it is set to <i>20</i> seconds.

10. Click on **OK** to complete the operation. The report opens and closes. The page is visible again, in the list **Network devices connection profiles** the new profile is listed as follows: *<profile-name> [PROTOCOL]*.

Editing Connection Profiles

At any time you can edit connection profiles.

Keep in mind that editing a connection profile associated with a device sets its status in the column **Versioning** to **Yes** on the page **All network devices**.

To edit a connection profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.
3. In the panel **Network devices connection profiles**, select in the list **Network devices connection profiles** the profile you want to edit.
4. Under the list, click on **EDIT**. The wizard **Edit a connection profile** opens.
5. Edit the profile. For more details, refer to the section [Adding Connection Profiles](#).
6. Click on **OK** to complete the operation. The report opens and closes. The page is visible again. If the profile you edited is used on a network device, its Versioning status is now **Yes**.

Deleting Connection Profiles

A connection profile can be deleted only if it is no longer used on any network device.

If you want to delete a connection profile already associated with a network device, you must edit the network device in question and select another connection profile. For more details, refer to the section [Configuring Versioning on One Device](#).

To delete a connection profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.
3. In the panel **Network devices connection profiles**, select in the list **Network devices connection profiles** the profile you want to delete.
4. Under the list, click on **DELETE**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The profile is no longer listed.

Configuring the Versioning of a Network Device

Configuring versioning implies to enable it and set associate your network device with a connection profile. It can be done on one network device at a time or on several network devices at once.

When versioning is enabled, you can manage the configuration files versioning from the pages **All configurations**, **Configuration** and **All network devices**.

Configuring Versioning on One Device

You can configure versioning for the network devices that support it directly on their properties page once you created a connection profile.

To configure versioning on a network device

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the panel **Configuration versioning properties**, click on **[EDIT]**. The wizard **Network device connection profile** opens.
4. Tick the box **Enable configuration versioning**. A drop-down list appears.
5. In the drop-down list **Network device connection profile**, select the connection profile of your choice.
6. Click on **[OK]** to complete the operation. The report opens and closes. In the panel, the line **Enable configuration versioning** is marked *yes* and the **Connection profile name** and **Connection protocol** of the device are displayed.

At any time you can change the connection profile associated with your network device. Follow the procedure above and select another connection profile.

Once you enabled versioning, the page **All configurations** allows to display the list of configuration file revisions and click on a revision number to open that version of the configuration file. To refresh a configuration files database, compare revisions or download revisions, refer to the dedicated sections.

Configuring Versioning on Several Devices

You can enable and set the configuration versioning on several network devices at once from the page **All network devices**. As the configuration is done from a single wizard, you have to make sure that the existing connection profile that you select matches the network devices' configuration needs.

To configure versioning on several network devices at once

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the network devices that you want to associate with your connection profile.
3. In the menu, **Tools > Set configuration versioning**. The wizard **Configuration versioning parameters** opens.
4. Tick the box **Enable configuration versioning**. A drop-down list appears.
5. In the drop-down list **Connection profile**, select the connection profile of your choice.
6. Click on **[OK]** to complete the operation. The report opens and closes. On the properties page of the selected devices, in the panel **Configuration versioning properties**: the line **Enable configuration versioning** is marked *yes* and the **Connection profile name** and **Connection protocol** of the device are displayed.

At any time you can change the connection profile associated with several network devices. Follow the procedure above and select another connection profile.

Once you enabled versioning, the page **All configurations** allows to display the list of configuration file revisions and click on a revision number to open that version of the configuration file. To refresh a configuration files database, compare revisions or download revisions, refer to the dedicated sections.

Refreshing a Configuration File

In addition to refreshing the database at any time, you can refresh the configuration files of the devices for which versioning is enabled. Refreshing the configuration file checks that the latest revision on the page **All configurations** matches the configuration of the network device you manage from the GUI at the time of the refresh.

Note that you can limit the refresh and only allow it for network devices which Status is *OK*. For more details, refer to the section [Limiting Configuration Refresh Based on the Network Device Status](#).

Refreshing Manually a Network Device Configuration File

The manual refresh allows to get the latest configuration available on a network device.

Keep in mind that the device parameter *SNMP transfer timeout* can impact the success of the refresh. For more details, refer to the section [Editing the SNMP Properties of a Network Device](#).

To manually refresh the configuration file of a device

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Tick the network device(s) you want to refresh.
3. In the menu, select  **Edit** > **Refresh**. The wizard **Refresh a network device** opens.
4. The box **Device data** is ticked. This box refreshes the database but not the configuration file of the selected network device(s). You can leave it ticked if you want. For more details, refer to the section [Refreshing the Network Devices Database](#).
5. Tick the box **Configuration versioning** to retrieve the latest revision of the configuration file of the selected device(s). This revision is retrieved only if any changes are detected.
6. Click on to complete the operation. The report opens and works for a while. The revision number is incremented if changes were detected. To compare the changes, go to the page **All configurations** of the network device and click on **Compare with previous version**.

Scheduling the Versioning Refresh on a Network Device

The configuration file refresh can be planned ahead and queried out automatically. The scheduled refresh can be set on any network device.

To set up a network device configuration file scheduled refresh

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Tick the network device(s) for which you want to schedule the refresh.
3. In the menu, select  **Edit** > **Scheduling** > **Configure refresh**. The wizard **Set refresh parameters** opens.
4. Configure the refresh frequency using the table below.

Table 63.5. Scheduled refresh parameters

Field	Description
Minute	Select the moment (o'clock, quarter past, half past or quarter to) or the frequency (in minutes) of the refresh.
Hour	Select a frequency (over the whole day or for a limited period of time each day), a set of hours or a specific hour per day for the refresh.
Date of the month	Select a specific day of the month or a frequency (every day) for the refresh.
Month	Select a specific month or a frequency (every month) for the refresh.
Day(s) of the week	Select a frequency (over the whole week or for a specific set of days) or a specific day of the week.

5. The box **Device data** is ticked. This box refreshes the database at the scheduled time but not the configuration file of the selected network device(s). You can leave it ticked if you want. For more details, refer to the section [Refreshing the Network Devices Database](#).
6. Tick the box **Configuration versioning** to retrieve the latest revision of the configuration file of the selected device(s). This revision is retrieved only if any changes are detected at the time of the scheduled refresh.
7. Click on to complete the operation. The report opens and closes. The page is visible again. The refresh frequency that you set is displayed in the panel **Refresh properties** on the properties page of your network device.

To disable a scheduled refresh, refer to the section [Disabling a Scheduled Refresh](#).

Comparing Configuration Files

From the page All configurations, you can access the page **Configuration** that provides a comparison of all the changes between two revision numbers, as raw data. The page Configuration in comparison display provides the following information:

- **Changed lines**, on the left-end side, indicates the number of lines edited between the two revisions.
- The buttons **Full diff** and **Short diff**, on the right-end side, allow to display the whole configuration files (Full diff), or only the lines that were edited (Short diff).

No matter the amount of information displayed, it is displayed as follows: on the left, the oldest revision where changes are highlighted in orange; on the right, the latest revision where changes are highlighted in green.

- **Previous changes**, on the right-end side, allows to compare older revisions of the configuration file, when possible. It actually compares the oldest revision with the previous revision of the configuration.
- **Next changes**, on the right-end side, allows you to navigate between comparisons of two revisions: toward more recent and/or older changes.
- , on the right-end side of each file revision allows to open the configuration file revision on its own, without the changes.

You can compare consecutive revisions of a device configuration file or even non-consecutive revisions of a configuration file or configuration files from different network devices.

Comparing Consecutive Revisions of a Configuration File

From the page *All configurations*, you can use the button **Compare with previous** to automatically open the page **Configuration** and display the revision of your choice compared with its previous version.

The very first revision available compares an empty file with that revision.

To compare consecutive revisions of a configuration file

1. In the sidebar, go to  **NetChange** > **Configurations**. The page **All configurations** opens.
2. In the column **Name**, click on the network device of your choice. Only the configuration file revisions of the device are listed; in the column **Revision**, the latest revision is named *r.<number> (last)*.
3. At the end of the line of the revision of your choice, click on **Compare with previous**. The page **Configuration** opens.

The oldest revision is on the left and the latest on the right. Each revision has its own gray area with the device revision details.

4. You can use the buttons to navigate the revisions changes, display the entire file with the changes or only the changes or even display only one revision. For more details, refer to the introduction of the section [Comparing Configuration Files](#).

Comparing Any Configuration File Revisions

From the page *All configurations*, you can tick several revisions - consecutive ones, the first and last revisions of a network device, two revisions belonging to different network devices - and compare them on the page **Configuration** using the menu.

To compare non-consecutive revisions of a configuration file

1. In the sidebar, go to  **NetChange** > **Configurations**. The page **All configurations** opens.
2. Filter the list of need be.
3. Tick the two revisions that you want to compare. It can be non-consecutive revisions of the same configuration file or from two different configuration files.
4. In the column **Name**, click on the network device of your choice. Only the configuration file revisions of the device are listed: in the column **Revision**, the latest revision is named *r.<number> (last)*.
5. At the end of the line of the revision of your choice, click on **Compare with previous**. The page **Configuration** opens.

The oldest revision is on the left and the latest on the right. Each revision has its own gray area with the device revision details.

6. You can use the buttons to navigate the revisions changes, display the entire file with the changes or only the changes or even display only one revision. For more details, refer to the introduction of the section [Comparing Configuration Files](#).

Monitoring the Configuration Versioning in the Logs

You can monitor versioning changes on the page **Syslog**, in the Administration module.

Versioning dedicated logs include: notifications if the device does not support it, if versioning is enabled/disabled, the new configurations, etc.

To display the versioning dedicated logs

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, make sure your local SOLIDserver is selected.
4. Under the menu, in the drop-down list **Services**, select *ipmsserver*.
5. In the column **Description**, filter the list using the keyword *configuration*.

The logs indicate when versioning has been enabled or disabled, they include the network device name and IP address.

Downloading Versioning Information

Once you enabled versioning, any configuration file in the database can be downloaded.

From the page **Configuration**, you can download one revision or the differences between two compared configuration files. From the page **All configurations**, you can display an archive file containing as many configuration file revisions as you need.

Downloading a Configuration File

You can download any configuration file revision displayed on the page **Configuration** in a *.txt* file.

To compare consecutive revisions of a configuration file

1. In the sidebar, go to  **NetChange** > **Configurations**. The page **All configurations** opens.
2. In the column **Revision**, click on the revision of your choice. The page **Configuration** opens.

Above the configuration file content, a gray area contains the network device details: its name, its IP address, and the date and time when the configuration version was saved.

3. On the right-end side of the network device details, click on **Download** and save the file. The downloaded file is named *<device-ip-address>_r<revision-number>.txt*.

Downloading Compared Configuration Files

You can download the compared display of two comparison files on the page **Configuration** in a *.diff* file.

Keep in mind that the downloaded file only contains the differences between the two revisions, even if you display the *Full diff*.

To download compared configuration file revisions

1. In the sidebar, go to  **NetChange** > **Configurations**. The page **All configurations** opens.

2. Compare the revisions of your choice. For more details, refer to the sections [Comparing Consecutive Revisions of a Configuration File](#) and [Comparing Any Configuration File Revisions](#).
3. Above the configuration files content, in the right-end side, click on **Download diff** and save the file. Only the differences between the files are downloaded.

Depending on the revisions you chose to compare, the downloaded file is named `<device-ip-address>_r<oldest-revision>_r<latest-revision>.diff` or `<device-1-ip>_r<revision-number>_<device-2-ip>_r<revision-number>.diff`.

Downloading Several Configuration Files in an Archive

You can download one or several revisions from the page **All configuration** in a `.tar` file.

To download configuration file revisions in an archive file

1. In the sidebar, go to **NetChange > Configurations**. The page **All configurations** opens.
2. Tick the revision(s) you want to download.
3. In the menu, select **Tools > Download device configurations**. The wizard **Download device configurations** opens.
4. Click on **OK** to complete the operation. The report opens, you can **DOWNLOAD** and save the file directly from the wizard.

The downloaded file is named `configs_<date>_<time>.tar` and is available in the module *Administration*, on the page **Local files listing**.

Configuring Advanced Versioning Options

Some registry database entries allow to customize the configuration files refresh possibilities and the passwords display within the files.

Configuring the Passwords Display in the Configuration Files

You can manage the versioning of your network devices, which is supported by most vendors, to automatically save all the changes in the configuration files. All the revisions of the files are then saved in SOLIDserver backup file.

Thanks to a registry database key, you can show or hide all or some of the passwords of the configuration file. By default, they are all hidden and we strongly recommend leaving it this way.

To configure the passwords display in the configuration files

Only users of the group `admin` can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in `module.iplocator.rancid.show_passwords` and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value listed. The **Registry database Edit a value** wizard opens.

5. In the field **Name**, the key name is displayed in a read-only gray field.
6. In the field **Value**, set the value that suits your needs:

Table 63.6. The accepted values of the configuration file password display key

Value	Description
0	Hide all the passwords in the configuration files. This is the default value.
1	Display only the encrypted passwords of the configuration files, in their encrypted form. All the non-encrypted passwords are hidden. Setting the key to 1 can be useful to keep track of the password changes without displaying them.
2	Display all the passwords of the configuration files. Setting the key to 2 is not recommended.

We strongly recommend leaving the default value 0.

7. Click on to complete the operation. The report opens and closes. The column **Value** contains the value you set.

Limiting Configuration Refresh Based on the Network Device Status

Thanks to a registry database key, you can prevent users from refreshing the configuration of the network devices with a Status different from OK. By default, the configuration refresh is allowed no matter the network device's status.

To set configuration refresh permissions based on status

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *module.iplocator.rancid.ignore_snmp_status* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value listed. The **Registry database Edit a value** wizard opens.
5. In the field **Name**, the key name is displayed in a read-only gray field.
6. In the field **Value**, set the value that suits your needs:

Table 63.7. The accepted values of the configuration file refresh key

Value	Description
0	Prevent users from refreshing the configuration of network devices which Status is not OK.
1	Allow configuration refresh of network devices no matter their Status. This is the default value.

7. Click on to complete the operation. The report opens and closes. The column **Value** contains the value you set.

Disabling Versioning on a Network Device

To disable versioning, you have to edit the network device network connection profile.

Keep in mind that **disabling the configuration files versioning on a device deletes all the configuration file revisions** saved in the database.

To disable versioning on a network device

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. At the end of the line of the network device of your choice, click on . The properties page opens.
3. In the panel **Configuration versioning properties**, click on **[EDIT]**. The wizard **Network device connection profile** opens.
4. Untick the box **Enable configuration versioning**. The page refreshes and the drop-down list **Network device connection profile** is no longer visible.
5. Click on **[OK]** to complete the operation. The report opens and closes. In the panel, only the line **Enable configuration versioning** remains, it is marked *no*.

Once disabled, the network device is marked *No* in the column **Versioning** on the page **All network devices**.

To disable versioning on several network devices at once

1. In the sidebar, go to  **NetChange > Network devices**. The page **All network devices** opens.
2. Tick the network devices that you want to associate with your connection profile.
3. In the menu, **Tools > Set configuration versioning**. The wizard **Configuration versioning parameters** opens.
4. Untick the box **Enable configuration versioning**. The page refreshes and the drop-down list **Connection profile** is no longer visible.
5. Click on **[OK]** to complete the operation. The report opens and closes. On the properties page of the selected devices, in the panel **Configuration versioning properties**: only the line **Enable configuration versioning** remains, it is marked *no*.

Once disabled, all the network devices selected are marked *No* in the column **Versioning** on the page **All network devices**.

Chapter 64. Managing Addresses

The page *All Addresses* lists the IP addresses configured for the interfaces of each network device managed from the page *All network devices*.

Browsing Addresses

The addresses are the second level of organization in NetChange, along with the routes, VLANs, ports and configurations.

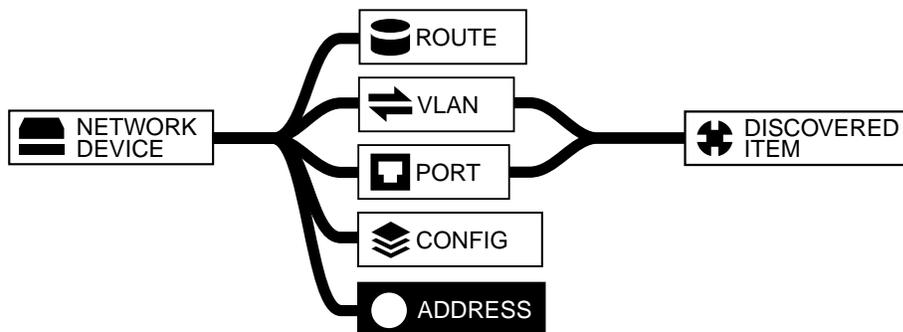


Figure 64.1. The Address in NetChange hierarchy

Browsing the Addresses Database

To display the list of configured IP addresses

1. In the sidebar, go to **NetChange > Addresses**. The page **All Addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To display the list of IP addresses configured for a specific network device, in the column **Network device**, click on the name of the device of your choice. The page refreshes.

Customizing the Display on the Page All Addresses

Users of the group *admin* can create customized column layouts. The button **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Table 64.1. Available columns on the page *All addresses*

Column	Description
IP address	The interface IP address.
Network address	The address of the network the IP address belongs to.
Interface name	The interface name.
Network device	The network device managing the interface the IP address is associated with.
VLAN number	The VLAN ID added to the network device.
VRF name	The name of the VRF as set on the network device.

Column	Description
VRF RD	The VRF Route Distinguisher. A unique 64-bits identifier, that can be composed of IP addresses or AS Numbers, that differentiates every set of routes on each VRF. For more details, refer to the part VRF .
Status	The interface status: <i>Active</i> or <i>Inactive</i> .

Keep in mind that you can use colored labels to differentiate at a glance IPv6 address containers. For more details, refer to the chapter [Managing IPv6 Labels](#).

Chapter 65. Managing Discovered Items

The discovered items are devices connected to the network devices, and usually edge devices (workstations, servers, printers...), identified through their MAC address. These devices are inserted in the database automatically after each discovery, and put in the history. This allows the administrator to know where a device (IP or MAC address) has been connected at a specific time and on which port of which device, in which VLAN, etc.

Browsing Discovered Items

The discovered items are the third and last level of the organization of the NetChange module.

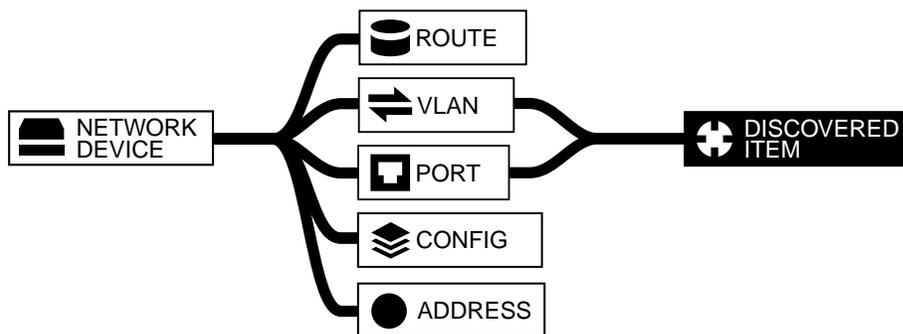


Figure 65.1. The discovered items in the NetChange hierarchy

Browsing the Discovered Items Database

To display the list of discovered items

1. In the sidebar, go to **NetChange** > **Discovered items**. The page **All discovered items** opens.
2. To display the list of discovered items of a specific network device, in the column **Network device**, click on the name of the device of your choice. The page refreshes.

One MAC address can be listed multiple times. It is listed for each unique IPv4 and/or IPv6 address associated with the discovered item.

The discovered items do not have a properties page, all the information is displayed on the page.

To display the discovered items history

1. In the sidebar, go to **NetChange** > **Discovered items**. The page **All discovered items** opens.
2. On the right-end side of the menu, click on **Discovered item history**. The page opens. For more details, refer to the section [Tracking the Discovered Items History of a Specific Device](#).

Customizing the Display on the Page All Discovered Items

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Some columns on the page provide specific information regarding the discovered items:

Table 65.1. Available columns on the page All discovered items

Column	Description
DNS name	The name of each discovered item is automatically retrieved by NetChange if it is declared in an A or PTR record in one of your DNS resolvers ^a .
Last seen	The time and date when the MAC address was last seen on a specific port and VLAN. This column is displayed by default on the page.
Discovered on	The time and date when the MAC address was seen for the first time on a specific port and VLAN. This column is displayed by default on the page.
First seen	The time and date when the MAC address was seen for the first time on a port, VLAN or network device. By default, this information is not retrieved, to retrieve it, refer to the section Enabling the Data Retrieval for the Column First Seen .
Source	The origin of the IP address, it indicates where it was retrieved.
	<i>Network device</i> : the IP address was retrieved either from the ARP table (for IPv4) or the NDP table (for IPv6) of the network device.
	<i>lease</i> : the IP address is a lease allocated in the DHCP.
	<i>static</i> : the IP address is a static reserved in the DHCP.
	<i>IPAM</i> : the IP address is assigned in the IPAM.
Network device IP	The IP address of the network device the item was discovered on.
IPv4 address	The IPv4 address of the discovered item, retrieved in the <i>Network device</i> ARP table. This column can be empty if the device has an IPv6 address or if no address was discovered.
IPv6 address	The IPv6 address of the discovered item, retrieved in the <i>Network device</i> NDP table. This column can be empty if the device has an IPv4 address or if no address was discovered. You can use colored labels to differentiate at a glance IPv6 address containers. For more details, refer to the chapter Managing IPv6 Labels .
VRF Name	The name of the VRF the MAC address belongs to.
VRF RD	The VRF Route Distinguisher.

^a[Setting the DNS Resolver](#) can be performed from the page Network configuration. NetChange does not retrieve the information in the records of the DNS servers you manage.

Enabling the Data Retrieval for the Column First Seen

The column *First seen* allows to display the time and date that the discovered items have ever been seen connected to a port, VLAN or network device. Retrieving that information can lower NetChange performances, which is why by default it is disabled.

To enable the retrieval of the data displayed in the column *First seen*, you must add a dedicated key in the registry database and reboot SOLIDserver.

To enable the data retrieval in the column First seen

1. Add the registry key

Only users of the group *admin* can perform this operation.

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- c. In the menu, click on **+ Add**. The wizard **Registry database Add an item** opens.
- d. In the field **Name**, type in `module.iplocator.activate_first_seen`.
- e. In the field **Value**, type in 1 to enable the data retrieval.
- f. Click on to complete the operation. The report opens and closes. The key is listed.

2. Reboot SOLIDserver

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Maintenance**, click on **Reboot SOLIDserver**. The wizard **Reboot the system** opens.
- c. Click on to complete the operation. You are logged out, the login page details the progression until the connection is lost.
- d. Refresh the page until the login page appears.

Once the reboot is complete and the login page is visible again, log in and go to the page *All discovered items*. In the column **First seen**, all the information available is displayed.

Refreshing the Discovered Items Database

You can refresh a selection of discovered items, whether they are associated with an IPv4 or IPv6 address from the page *All discovered items*.

To refresh manually or schedule the refresh of a network device and all its discovered items, refer to the section [Refreshing the Network Devices Database](#).

To refresh the discovered items manually

1. In the sidebar, go to  **NetChange** > **Discovered items**. The page **All discovered items** opens.
2. Tick the discovered item(s) you want to refresh.
3. In the menu, select  **Edit** > **Refresh**. The wizard **Refresh a network device** opens.
4. Click on to complete the operation. The report opens and closes. The page is visible again.

Populating Device Manager

The same way you can create a discovered network device into Device Manager, you can populate Device Manager with a selection of discovered items. This allows to create the corresponding device and interface in Device Manager, with the MAC address differentiating each interface.

For more details, refer to the section [Automatically Adding Discovered Items in Device Manager](#).

Creating the IP Address of a Discovered Item in the IPAM

You can assign IPv4 and IPv6 addresses in the IPAM from the page *All discovered items*.

To successfully create IP addresses in the IPAM from NetChange, make sure that:

1. The IP address of the discovered item is actually retrieved and displayed. It might not be available if the router has not found any equivalence between the MAC and IP address, neither from the equipment nor from the DHCP.
2. The IP address of the discovered item is available in the IPAM. The option allows to:
 - Select a *Target space*, either the same space as the one set for the network device managing your item or another space, as long as this space contains a terminal network with free addresses that you can assign to your discovered items. For more details, regarding the space selection at network level, refer to the section [Adding Network Devices](#).
 - Tick the box *Use best space* to automatically find a space that can receive the discovered item in the smallest terminal network available.

To create the IP address of a discovered item in the IPAM

1. In the sidebar, go to  **NetChange** > **Discovered items**. The page **All discovered items** opens.
2. Tick the discovered item(s) you want to create in the IPAM.
3. In the menu, select  **Tools** > **Create IP address in the IPAM**. The wizard **Create IP addresses in the IPAM** opens.
4. Specify the space where the IP address of the items is created in the IPAM:
 - a. Either tick the box **Use best space** to place the discovered item(s) into the space containing the smallest network able to receive them.
 - b. Or in the drop-down list **Target space**, select the space that suits your needs.
5. Click on to complete the operation. The report opens and closes and the list is visible again. The IP address is listed on the page **All addresses**, in the IPAM.

Purging the Discovered Items Database

At any time you can purge the discovered items database of all your network devices at once, that is to say remove MAC addresses older than a certain period from the page *All discovered items*.

All MAC addresses discovered prior to the period you set are removed from the page and only available on the page *Discovered items history*. For more details, refer to the section [Tracking the Discovered Items History of a Specific Device](#).

Note that you can also configure or refine the purge frequency via the rule 008, for more details refer to the section [Keeping NetChange Data Up-to-date](#).

To purge the discovered items database of all the network devices

1. In the sidebar, go to  **NetChange** > **Discovered items**. The page **All discovered items** opens.

2. In the menu, select **Tools > Administration > Purge MAC addresses**. The wizard **Purge MAC addresses history** opens.
3. In the field **Number of hours to keep**, type in the value of your choice.
4. Click on **OK** to complete the operation. The wizard closes. Only the items which MAC address was discovered during the period you set are listed.

Tracking the Discovered Items History of a Specific Device

NetChange database is not erased automatically. However, from the page *Discovered items history* you can have an overview of the movements of a specific edge device during previous discoveries.

That way, using a MAC address, you can see the different IP addresses an edge device had, at different periods of time, which switch and port it was connected and which VLAN it belonged to. This function also allows to track laptops on the network and see on which switches and ports they have been successively connected to.

To manage the columns display, refer to the section [Customizing the Display on the Page All Discovered Items](#).

To display the Discovered items history of a specific network device

1. In the sidebar, go to  **NetChange > Discovered items**. The page **All discovered items** opens.
2. In the column **Network device**, click on the name of the network device of your choice. The page refreshes, the device name is displayed in the breadcrumb.
3. On the right-end side of the menu, click on  **Discovered item history**. The page opens and displays only the history of the network devices of the network you chose.
4. If need be, filter the list using the search engine of the column(s) of your choice.

To limit the number of discovered items saved in the history, you can configure and enable a rule dedicated to purging the NetChange database, including discovered items. For more details, refer to the section [Keeping NetChange Data Up-to-date](#).

Chapter 66. Managing Statistics

NetChange can provide a set of specific statistics. These statistics are all displayed as pie and bar charts that can present vendors, speed, usage, etc.

Displaying NetChange Statistics

By default, NetChange dashboard displays four gadgets (pie charts) that present the database network devices vendors, number of ports per device, ports status and ports speed.

To display NetChange gadgets

1. From any page, in the top bar, select  **My account** > **My Gadgets**. The page **My Gadgets** opens.
2. In the breadcrumb, click on **Gadgets Library**. The page **Gadgets Library** opens.
3. In the list, tick the gadget(s) of your choice: *Number of NetChange ports per device*, *NetChange network devices vendor*, *NetChange active ports speed (bps)* and/or *NetChange ports status*.
4. In the menu, select  **Edit** > **Assign Gadget(s)**. The gadget configuration wizard opens.
5. In the list **Available**, select a module and click on .
6. The module name is moved to the list **Configured**.
7. Click on  to complete the operation. The report opens and closes. The gadgets are displayed on the selected dashboards and listed in the column **Dashboard**.

In addition to these holistic charts, NetChange provides specific charts on the network devices and port properties pages.

Displaying Network Device Statistics

The properties page of the network devices contains a chart summing up the device ports status.

To display the statistics of a network device

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. At the end of the line of the device of your choice, click on . The properties page opens.
3. Open the panel **Network device ports status** to display the graph.

Displaying Port Statistics

The ports properties page provides a set of charts that display accurate data if you enable the rule 067.

Enabling the Rule That Retrieves Ports Information

The rule 067 allows to retrieve ports information and display it on the ports properties page. By default, it is listed among the *Rules* but disabled, so you have to activate it to retrieve data.

To activate the collection of data rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. Through the columns filters, look for the rule 067 called *Charts of NetChange ports bandwidth usage*.
4. Tick the rule.
5. In the menu, select  **Edit** > **Enable**. The wizard **Enable** opens.
6. Click on to complete the operation.

Displaying a Port Charts

On the port properties page, you can display four different charts: *NetChange port traffic*, *NetChange port broadcast traffic*, *NetChange port unicast traffic* and *NetChange port error packet traffic*.

To display the statistics of a port

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Click on the name of the device of your choice. The page **All ports** of the device opens.
3. At the end of the line of the port of your choice, click on . The properties page opens.
4. Click on  in the upper right corner of the page to open all the panels.

The charts contain In and Out parameters. To better understand them, refer to the table below.

Table 66.1. NetChange ports charts information

Information	Description
In octets	The total number of octets received on the port, including framing characters.
Out octets	The total number of octets transmitted out of the port, including framing characters.
In broadcast	The number of non-unicast (i.e., subnetwork- broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol.
Out broadcast	The total number of packets that higher-level protocols requested be transmitted to a non- unicast address, including those that were discarded or not sent.
In unicast	The number of unicast packets delivered to a higher-layer protocol.
Out unicast	The total number of packets that higher-level protocols requested be transmitted to a subnetwork- unicast address, including those that were discarded or not sent.

Chapter 67. Monitoring and Tuning

NetChange offers a set of options to monitor and tune your network devices and discovered items.

Generating NetChange Reports

EfficientIP provides NetChange dedicated reports for network devices. The reports on inconsistencies might be empty if the devices configuration is correct.

Table 67.1. Available NetChange reports

Page	Report
All network devices	Network Devices Properties
	NetChange/IPAM/DHCP data comparison
	Network devices summary

For more details regarding the reports and their generation, refer to the section [Managing Reports](#).

Keeping NetChange Data Up-to-date

To always have an up-to-date data, you should periodically refresh the network devices using scheduling, as explained in the chapter [Managing Network Devices](#). You should also remove old data from the history to speed up the processes and have only the relevant IP or MAC address information. The choice of periodicity depends completely on your environment and what you intend to do with NetChange: you may need to have a history of all movements (so you might need to purge the database every month or trimester), or you may need only the most relevant data when looking for a host (so you might want to purge every week).

To configure the purge frequency of the data listed on the page *All discovered items*, you should edit the rule 008.

To configure and enable the rule 008 that purges NetChange History

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #** search field, type in *008*. The rule *Purge NetChange history* is listed.
4. **Configure the rule**
 - a. In the column **Instance**, click on *iplocator_cron_purge_history*. The rule properties page opens.
 - b. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a rule** opens.
 - c. Click on **[NEXT]**. The page **Rule filters** appears.
 - d. If you want to schedule the purge, configure the fields according to the table below:

Table 67.2. Rule filters parameters

Field	Description
Day(s) of the week	Select a day or a period of days.
Date of the month	Select a date.

Field	Description
Month	Select a month.
Hour	Select a specific time or one of the available schedules.
Minute	Select a period of time, minutes-wise.

- e. Click on **NEXT**. The page **Rule parameters** appears.

Table 67.3. Rule parameters

Field	Description
Purge MAC/port history	In this drop-down list, you can keep (<i>No</i>) or delete (<i>Yes</i>) the history data of all the MAC addresses that were discovered through a port, when their lifespan exceeds the <i>Number of days to keep</i> . By default, <i>Yes</i> is selected.
Purge MAC/IP history	In this drop-down list, you can keep (<i>No</i>) or delete (<i>Yes</i>) the association history of all the MAC addresses that were retrieved from the ARP table of the network device, when their lifespan exceeds the <i>Number of days to keep</i> . By default, <i>Yes</i> is selected.
Number of days to keep	In this field, type in the maximum number of days that you want to keep NetChange history, this includes the discovered item and time entries are listed in the History view.

- f. Click on **OK** to complete the operation. The report opens and closes. The rule properties page is visible again.

5. Enable the rule

- Tick the rule *008* and, in the menu, select **Edit > Enable**. The wizard **Enable** opens.
- Click on **OK** to complete the operation. The wizard closes, the page refreshes. The rule is **Enabled**.

Synchronizing the Network Devices with a CSV File

Thanks to the rule *010*, it is possible to automatically synchronize the list of network devices with a local file containing on each line the IP address of a network device to be managed in NetChange. If the CSV file contains devices that are not managed by NetChange yet, they are automatically imported during the synchronization.

To activate the automated retrieval of devices from a specific CSV import

- In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Expert**, click on **Rules**. The page **Rules** opens.
- In the column **Name**, type in the rule name *Synchronize NetChange network devices with a CSV file*.
- At the end of the line, click on **⊞**. The properties page opens.
- In the panel **Main properties**, click on **EDIT**. The wizard **Edit a rule** opens.
- Click on **NEXT**. The page **Rule filters** opens.
- Define the frequency of execution of the rule. The default frequency is every 5 minutes.
- Click on **NEXT**. The page **Rule parameters** opens.

9. Fill in the following fields:

Table 67.4. Rule parameters

Field	Description
Local CSV file	Type in the path toward the CSV file containing the list of addresses. This CSV file must already be part of the local SOLIDserver file system.
Devices missing in CSV file	You can <i>Delete</i> or perform <i>No action</i> over the devices already listed in the page <i>All network devices</i> but not present in the CSV file to import.
Site id	Select the space where you want the IP address of the devices to be imported. This field is required.
Delimiter	Select the delimiter that separates the data in your CSV file: It can be a comma, a semicolon or a tab. By default, the comma is selected. This field is optional.

10. Click on to complete the operation.

Once the rule is configured, you have to enable it. Before following the procedure below, check in the column **Status** if the rule is marked *Disabled* or *OK* (i.e. It is enabled). Note that if you enable the rule before configuring it, no action is performed for lack of specifications (without a path toward a CSV file the options were configured with no file to perform them on).

To enable the devices synchronization rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Name**, type in the rule name *Synchronize NetChange network devices with a CSV file*.
4. Tick the rule.
5. In the menu, select  **Edit** > **Enable**. The wizard **Enable** opens.
6. Click on to complete the operation. The report opens and closes. The page is visible again and the rule is marked  **OK** in the column **Status**.

Customizing the Type of Network Devices

On the page *All network devices*, the column **Type** displays partial vendor and/or model information about each network device. You can customize the content of this column via CLI.

By default, the file *custom_sysoid.csv* allows to include the system OID of the network device of your choice and the name you want to display in the column *Type*. Keep in mind that:

- Only users with administrative rights over SSH connections to SOLIDserver can edit the file.
- On the page *All network devices*, the OID of each network devices is available in the column **System OID**.
- You must refresh the devices to display the new information in the column.

To customize the type of a network device

1. Edit the file via CLI:
 - a. Connect to your appliance using an SSH session or a port console and *root* credentials.

- b. Edit the file `/data1/etc/custom_sysoid.csv` to include a line for each network device.

The system OID of each device and the name to display must be separated by a comma.

```
<Full.system.OID>,<Information displayed in the column>
#For instance, renaming EfficientIP network devices could look as follows in the file:
#.1.3.6.1.4.1.2440,EfficientIP network device
```

2. Refresh the network devices:

- a. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
- b. Tick the network device(s) you want to refresh.
- c. In the menu, select  **Edit** > **Refresh**. The wizard **Refresh a network device** opens.
- d. Tick the **Device data** to refresh the database of the selected network device(s).
- e. Click on **OK** to complete the operation. The report opens and works for a while.

When the refresh is over, a report might appear. You can download this report in the format of your choice.

- f. Click on **CLOSE** to go back to the page **All network devices**. The page refreshes. In the column **Type**, the information regarding the device(s) matching the OID you specified in the file is updated.

Part XII. Workflow

The Workflow is a request-based module allowing standard users to ask for changes in the IPAM and DNS:

- DNS **zones** addition, edition or deletion.
- IPAM **non-terminal networks** addition, edition or deletion. These can be block-type or subnet-type networks.
- IPAM **terminal networks** addition, edition or deletion.
- IPAM **pools** addition, edition or deletion.
- IPAM **IP addresses** addition, edition or deletion.

All the requests are listed on both pages of the module:

- **Outgoing requests**, where users - requestors - can create requests and monitor their status evolution. For more details, refer to the chapter [Managing Outgoing Requests](#).
- **Incoming requests**: where request managers and administrators can deal with the user requests. For more details, refer to the chapter [Managing Incoming Requests](#).

In addition, you can:

- **Automate the requests' execution**. You can set up automated execution of incoming requests, either using the Workflow default classes and the button *Execute* or include pending operations to the class objects of their customized classes. For more details, refer to the chapter [Executing Requests](#).
- **Customize the requests acceptance cycle**. Some registry database entries determine the requests' life cycle using their statuses. For more details, refer to the chapter [Customizing the Requests Administration](#).

Keep in mind that to use the Workflow at the best of its potential you must:

1. **Grant sufficient rights to requestors and request managers**: the group they belong to needs to be granted the appropriate IPAM, DNS and/or Workflow rights. For more details, refer to the section [Managing the Rights of a Group of Users](#) in the chapter [Managing Groups](#).
2. **Grant users access to request classes**, the existing classes or the ones you created. For more details, refer to the chapter [Granting Access to Workflow Classes](#).
3. **Customize the page Incoming requests if need be**. For more details, refer to the chapter [Customizing the Requests Administration](#).
4. **Grant relevant users access to the Workflow pages**, that way they can create or deal with the requests.
5. **Executing the action required in the requests** if they are accepted.

Note that from the module **Dashboards**, you can gather gadgets and charts on *Workflow dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 68. Granting Access to Workflow Classes

As every request is based on a specific Workflow class, users need to be granted access to the relevant ones. That way, they can select a class when adding a request and fill in the fields defined through the class.

There are five classes dedicated to Workflow requests. They define all the mandatory fields when asking for the addition, edition or deletion of the object they are named after:

- *request_dns_zone* is dedicated to requests regarding DNS zones.
- *request_ip_block* is dedicated to requests regarding IPv4 non-terminal networks, block-type or subnet-type.
- *request_ip_subnet* is dedicated to requests regarding IPv4 terminal networks.
- *request_ip_pool* is dedicated to requests regarding IPv4 pools.
- *request_ip_address* is dedicated to requests regarding IPv4 addresses.

The users that do not have access to Workflow request classes are not able to properly complete the request addition wizard: the request addition wizard is still available, but it is impossible to define the needed containers or resources to apply the requested changes to.

Obviously, you can add your own Workflow request classes. These classes must be dedicated to the Module Workflow and the Type Request. For more details, refer to the chapter [Configuring Classes](#).

Keep in mind that in this case, the *Execute* option is not available in the page Incoming requests. For more details, refer to the section [Executing Requests Using the Execute Option](#).

To add an existing request class as group resource

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Click on the name of the group of your choice¹. The page **Resources** opens.
4. In the menu, select **+ Add > Resources > Classes**. The wizard **Administration: Classes** opens.
5. In the search engine of the column **Type**, type in *request* to filter the list.
6. Tick the class(es) you want to grant the group access to. Keep in mind that the default Workflow classes are *request_dns_zone*, *request_ip_block*, *request_ip_subnet*, *request_ip_pool* and *request_ip_address*.
7. Click on to complete the operation. The report opens and closes. The page refreshes, the list of resources now includes the selected class(es).

Once the classes of your choice are part of the resources of a group, its users can choose from one of them when requesting the addition, edition or deletion of objects in the DNS or IPAM database.

¹Any group EXCEPT the *admin* group as, by default, it has authority over all the resources of SOLIDserver database.

Chapter 69. Managing Outgoing Requests

From the page Outgoing requests, users with sufficient Workflow rights can:

- Add requests for addition/edition/deletion.
- Edit the requests they created.
- Cancel the requests they created.

The requests management respects the groups hierarchy by default. Therefore, once created if the user belongs to a group that has a parent group, then by default the request can be dealt with by all the users of the parent group as well as the users of the group *admin*. If the users want the request to be dealt with by specific users, they can set a managing group when creating or editing the request.

Browsing Outgoing Requests

The Outgoing requests page is one of the two pages of the module. Requestors use this page to add, edit and cancel requests.

Browsing the Outgoing Requests Database

To display the list of outgoing requests

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. You can filter the list using the column search engines.

To display an outgoing request properties page

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. To display a request properties page you can:
 - a. Click on the name of the request of your choice. The properties page opens.
 - b. At the end of the line of the incoming request of your choice, click on . The properties page opens.
3. Click on  to expand all the panels.

Customizing the Display on the Page Outgoing Requests

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding Requests for Creation

Users with sufficient rights can add requests asking for objects creation. This request can only contain basic information regarding the object, the advanced properties can only be configured by administrators and request managers when they execute the request.

The procedure below uses the *request_dns_zone* class as an example.

Reminder

To add a DNS zone creation request, the group of the user must have at least been granted the following **rights**:

- In the panel Workflow, all the rights that suit your needs
- In the panel DNS, the right *Display: DNS servers list*.

To add a DNS zone creation request, the group of the user must include among its **resources**:

- At least one server to grant access to all the objects it contains.

To add a request for creation

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. In the menu, click on **+ Add**. The wizard **Add a request** opens.
3. In the list **Workflow request class**, select *request_dns_zone* .
4. Click on **NEXT**. The page **Requesting: Zone** opens.
5. In the drop-down list **Action requested**, select *New (Create)*. The page refreshes.
6. In the drop-down list **DNS Server**, select the server of your choice.
7. If your server contains views, in the drop-down list **DNS view**, select the view of your choice.
8. In the field **DNS zone**, name your zone.
9. In the drop-down list **Zone type**, you can select either *master*, *slave*, *stub* or *forward*. By default, *master* is selected.
10. In the field **Motivation**, type in a text or a maximum of 3000 characters explaining the reason for the zone addition request.
11. Click on **NEXT**. The last page of the wizard opens.
12. If you want, you can select a group to manage your request as described in the table below. Otherwise, users of the group *admin* can manage it.

Table 69.1. Group of users administrating the request dedicated fields

Parameter	Description
Set a group to manage the request	Allows to decide that a specific group of users manages the request. By default, it is set to <i>No</i> .
No	Only the users from the requesting user parent group or the users of the group <i>admin</i> can manage your request.
Yes	You can select an existing group of users to manage your request. Once the option is selected, the drop-down list <i>Managing group</i> appears.
Managing group	Select an existing group. This action allows the request managers of the group to accept the request and deal with it, or deny it.

- Click on **OK** to complete the operation. The report opens and closes. The request is listed and marked as *New* in the column **Status**.

Each request is named as follows: *<request-number>-<requestor>*, where *<request-number>* corresponds to the number of requests in the Workflow database and not the number of requests of a particular requestor.

On the request properties page, the **Main properties** and **Request parameters** sum up the request details.

Adding Requests for Edition

Users with sufficient rights can add requests asking for objects edition. This request can only contain basic information regarding the object, the advanced properties can only be configured by administrators and request managers when they execute the request.

The edition request only applies to the values that you can usually edit in each module, so:

- You cannot ask for the edition of anything configured for DNS zones.
- You can only ask for the edition of the name of the networks, pools and addresses.

The procedure below uses the default *request_ip_address* class as an example.

Reminder

To add an IPv4 address edition request, the group of the user must have at least been granted the following **rights**:

- In the panel Workflow, all the rights that suit your needs
- In the panel IPAM, the right *Display: spaces list*.

To add an IPv4 address edition request, the group of the user must include among its **resources**:

- At least one space, which grants access to all the objects it contains.

To add a request for edition

- In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
- In the menu, click on **+ Add**. The wizard **Add a request** opens.
- In the list **Workflow request class**, select *request_ip_address*.
- Click on **NEXT**. The page **Requesting: IP address** opens.
- In the drop-down list **Action requested**, select *Modify*. The page refreshes.
- Click on **NEXT**. The next page opens.
- In the list **Choose a terminal network**, select the terminal network containing the IP address you want to edit. Once selected, the line is highlighted in blue.
- Click on **NEXT**. The next page of the wizard opens.
- In the field **Name to edit**, type in the first letter(s) of the IP name. The auto-completion provides a list of addresses matching these letters, select the one you want to edit.
- In the field **IP address name**, type in the new name of the address.
- In the gray field **IP address**, the IP address is displayed as a reminder.

12. In the field **Motivation**, type in a text or a maximum of 3000 characters explaining the reason for the IP address edition request.
13. Click on **[NEXT]**. The last page of the wizard opens.
14. If you want, you can select a group to manage your request as described in the table below. Otherwise, users of the group *admin* can manage it.

Table 69.2. Group of users administrating the request dedicated fields

Parameter	Description
Set a group to manage the request	Allows to decide that a specific group of users manages the request. By default, it is set to <i>No</i> .
No	Only the users from the requesting user parent group or the users of the group <i>admin</i> can manage your request.
Yes	You can select an existing group of users to manage your request. Once the option is selected, the drop-down list <i>Managing group</i> appears.
Managing group	Select an existing group. This action allows the request managers of the group to accept the request and deal with it, or deny it.

15. Click on **[OK]** to complete the operation. The report opens and closes. The request is listed. It is marked *New* in the column **Status** and *Modified* in the column **Action**.

Each request is named as follows: *<request-number>-<requestor>*, where *<request-number>* corresponds to the number of requests in the *Workflow* database and not the number of requests of a particular requestor.

On the request properties page, the **Main properties** and **Request parameters** sum up the request details.

Adding Requests for Deletion

Users with sufficient rights can add requests asking for objects deletion. Besides, the terminal networks deletion request can be created from the page *All networks*. For more details, refer to the section [Creating a Network Deletion Request From the IPAM](#).

Creating a Request for Deletion

In the procedure below we use the default *request_ip_subnet* class as an example.

Note that to the class *request_ip_block* applies to non-terminal networks of both subnet and block type.

Reminder

To create a terminal network deletion request, the group of the user must have at least been granted the following **rights**:

- In the panel *Workflow*, all the rights that suit your needs
- In the panel *IPAM*, the right *Display: spaces list*.

To edit an IP address related request, the group of the user must include among its **resources**:

- At least one space to grant access to all the objects it contains.

To add a request for deletion

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. In the menu, click on **+ Add**. The wizard **Add a request** opens.
3. In the list **Workflow request class**, select *request_ip_subnet*.
4. Click on **[NEXT]**. The page **Requesting: Network (subnet)** opens.
5. In the drop-down list **Action requested**, select *Delete*. The page refreshes.
6. Click on **[NEXT]**. The next page opens.
7. In the list **Choose a terminal network**, select the terminal network you want to delete. Once selected, the line is highlighted in blue.
8. Click on **[NEXT]**. The next page of the wizard opens.
9. The fields **Network name**, **Network address**, **Netmask**, **Prefix** and **Comments** display the selected network information as a reminder.
10. In the field **Motivation**, type in a text or a maximum of 3000 characters explaining the reason for the network deletion request.
11. Click on **[NEXT]**. The last page of the wizard opens.
12. If you want, you can select a group to manage your request as described in the table below. Otherwise, users of the group *admin* can manage it.

Table 69.3. Group of users administrating the request dedicated fields

Parameter	Description
Set a group to manage the request	Allows to decide that a specific group of users manages the request. By default, it is set to <i>No</i> .
No	Only the users from the requesting user parent group or the users of the group <i>admin</i> can manage your request.
Yes	You can select an existing group of users to manage your request. Once the option is selected, the drop-down list <i>Managing group</i> appears.
Managing group	Select an existing group. This action allows the request managers of the group to accept the request and deal with it, or deny it.

13. Click on **[OK]** to complete the operation. The report opens and closes. The request is listed. It is marked *New* in the column **Status** and *Delete* in the column **Action**.

Each request is named as follows: *<request-number>-<requestor>*, where *<request-number>* corresponds to the number of requests in the *Workflow* database and not the number of requests of a particular requestor.

On the request properties page, the **Main properties** and **Request parameters** sum up the request details.

Creating a Network Deletion Request From the IPAM

In addition to the standard addition of requests for deletion from the page *Outgoing request*, it is possible to create a request for network deletion from the page *All networks*.

It automatically creates a request in the *Workflow* using the default class *request_ip_subnet* for terminal networks or *request_ip_block* for non-terminal networks (block-type or subnet-type).

Reminder

To add a network deletion request, the group of the user must have at least been granted the following **rights**:

- In the panel **Workflow**, all the rights that suit your needs.
- In the panel **IPAM**, the right *Display: spaces list*.

To edit an IP address related request, the group of the user must include among its **resources**:

- At least one space to grant access to all the objects it contains.

To add a request for networks deletion from the IPAM

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. Filter the list if need be.
3. Tick the network for which you want to create a request for deletion.
4. In the menu, select **Edit > Request for deletion**. The wizard **Workflow Creating a network deletion request** opens.
5. In the drop-down list **Source**, select a group if you belong to several groups. If you only belong to one group, it is automatically selected. The managing group of the selected group handles the request.
6. Click on **OK** to complete the operation. The report opens and closes.
7. In the sidebar, go to **Workflow > Outgoing requests**. The request is listed. It is marked *New* in the column **Status** and *Delete* in the column **Action**.

Editing Requests

Once you created a request, you can edit its details or provide additional information via a note and/or file upload.

Editing a Request Details

Once created, the details of a request can be edited to a certain point: you cannot edit the action required or the object it applies to.

To edit a request from the outgoing request page

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on **Edit**. The wizard opens.
4. Edit the fields as needed. Only the fields with a white background can be edited.
5. Once you get to the last page of the wizard, click on **OK** to complete the operation. The report opens and closes. The changes are visible on the properties page panel **Request parameters**.

To edit a request from its properties page

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on **Properties**. The request properties page opens.

4. In the panel **Main properties**, click on **EDIT**. The wizard opens.
5. Edit the fields as needed. Only the fields with a white background can be edited.
6. Once you get to the last page of the wizard, click on **OK** to complete the operation. The report opens and closes. The changes are visible in the panel **Request parameters**.

Adding Information to a Request

From a request properties page, users can add information to a request via notes and files upload.

Uploading a File

Requestors can add up to 10 files to their request. They cannot upload more than 5 Mb of files.

To upload a file to a request properties page

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on **⊞**. The request properties page opens.
4. In the panel **Upload file**, click on **EDIT**. The wizard **Upload file** opens.
5. Click on **BROWSE** to select the file of your choice on your local computer.
6. Once selected, it is displayed in the fields **File name** and **Final value**.
7. Click on to **⊞** add the file to the list **Attached files**. Repeat these actions for as many files as you want.
8. Click on **OK** to complete the operation. The report opens and closes. The panel **Upload file** contains the file(s).

Adding a note

Requestors can add notes to their request in addition to the *Motivation* expressed when creating the request.

To add a note to a request

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on **⊞**. The request properties page opens.
4. In the panel **Note**, click on **EDIT**. The wizard **Enter a note** opens.
5. Click on **ADD**. In the field **List**, the line *new_<number>* appears.
6. In the field **Note**, type in your note. The note must not include special characters or exceed 3993 characters.
7. Click on **ADD** to save it. The note is saved. In the field **List**, the note is now displayed as follows: *<date> <time> <beginning-of-note> [author]*. Repeat these actions for as many notes as needed.
8. Click on **OK** to complete the operation. The report opens and closes. The panel **Note** displays the note(s).

Canceling Requests

At any time, you can cancel a request you created. By default, this action is only possible if the request status is New. Once it is handled or accepted, you can no longer cancel it.

Once canceled, you no longer see it on the page, only request managers can still see it.

To cancel a request

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Tick the request you want to cancel.
3. In the menu, select **Edit > Cancel**. The wizard **Status edition** opens.
4. In the field **Enter a note**, you can type in a text of 3000 characters at most, explaining why you want the request canceled.
5. Click on **OK** to complete the operation. The report opens and displays the cancellation report status.
6. You can click on **CSV (DATA)**, **TEXT**, **HTML** or **EXCEL** to download the cancellation report in the corresponding format.
7. Click on **CLOSE** to close the wizard. The page refreshes. The request is no longer listed on the page.

Chapter 70. Managing Incoming Requests

From the page Incoming requests, administrators or *request managers* can:

1. Deal with pending requests using the default options of the menu *Edit*: handle, edit, execute, reject, finish and finally delete the requests.
2. Deal with pending requests using custom options. The available options would then depend on the administrator configuration and intern use of the module.

Keep in mind that by *request managers*, we mean users belonging to a group with sufficient rights and resources. Make sure they belong to a group configured with:

- All the Workflow rights, to be able to manage the requests completely.
- All the DNS and IPAM objects that regular users can create requests for among the group resources.
- All the relevant IPAM and DNS rights that allow them to comply with the request.

Browsing Incoming Requests

The Incoming requests page allows the request managers and administrators to deal with the requests.

Browsing the Incoming Requests Database

To display the list of incoming requests

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. You can filter the list using the column search engines.

To display an incoming request properties page

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. At the end of the line of the incoming request of your choice, click on . The properties page opens.
3. Click on  to expand all the panels.

Customizing the Display on the Page Incoming Requests

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Managing the Requests Content

Once a request is created, it is listed on both pages of the module. Administrators and request managers deal with them from the Incoming requests.

On the request properties page are displayed all the request details as well as the requestor notes and uploaded files. In the Request history are listed all the administrators and request managers notes added when editing the request status.

To download an uploaded file

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on . The request properties page opens.
4. Click on  to expand all the panels.
5. In the panel **Upload file**, click on the name of the file you want to download.

To display notes

1. In the sidebar, go to **Workflow > Outgoing requests**. The page **Outgoing requests** opens.
2. Right-click over the **Name** of the request you want to edit. The contextual menu opens.
3. Click on . The request properties page opens.
4. In the panel **Note**, all the notes are displayed under the Date and User. You can scroll down if there are several notes.

The administrators and request managers can also add notes and upload files. For more details, refer to the section [Adding Information to a Request](#).

Administrating Requests Using the Default Statuses and Options

There are 6 default statuses on the page Incoming requests of the Workflow module. They allow administrators and request managers to see what requests need to be dealt with. All these statuses can be set from the menu *Edit* on the page Incoming requests. Except *canceled* that has to be set from the page Outgoing requests.

Every time a request status is edited, it sends an email to the user who requested it to inform them of the request evolution. Therefore, make sure your requesting users profile is set up properly. For more details, refer to the chapter [Managing Users](#).

Only the *Archive* option does not correspond to any status as it basically deletes the request from the page and stores it on the page *Local files listing*.

Table 70.1. Request statuses

Status	Description
New	The request was created on the page Outgoing requests and has not been dealt with yet.
Handled	The request was acknowledged by a request manager or administrator, it still has to be accepted or rejected. This status can only be set from the page Incoming requests.
Accepted	The request was accepted by the request manager or administrator handling it: the requested creation/edition/deletion will be performed. This status can only be set from the page Incoming requests.
Rejected	The request was denied by the request manager or administrator handling it. Whatever was requested is ignored. This status can only be set from the page Incoming requests.

Status	Description
Canceled	The request was canceled by the requestor. It is no longer requested creation/edition/deletion has to be ignored by the request manager or administrator handling it. This status can only be set from the page Outgoing requests.
Finished	The requested creation/edition/deletion was performed. This status can only be set from the page Incoming requests.

By default, the requests managers can set these statuses as long as they respect the following:

- *New* requests can be handled, accepted, rejected and canceled.
- *Handled* requests can be accepted, rejected and canceled.
- *Accepted* requests can only be dealt with and finished.
- *Rejected* requests can only be archived.
- *Canceled* requests can only be archived.
- *Finished* requests can only be archived.

Using the default options and statuses is useful as it allows to use the *Execute* option. This option allows to execute a request from the Incoming requests directly. For more details, refer to the section [Executing Requests Using the Execute Option](#).

Handling Requests

The request managers and administrators can at any point handle *New* requests.

You cannot handle *Accepted*, *Rejected* or *Finished* requests.

To handle a request

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. Tick the request(s) you want to handle.
3. In the menu, select **Edit > Handle**. The wizard **Status edition** opens.
4. In the field **Enter a note**, you can type in a reason for accepting or the user performing the task. This text is available on the request properties page, in the **Request history** panel.
5. Click on **OK** to complete the operation. The report opens and indicates the operation success.
6. Click on **CLOSE** to go back to the page Incoming requests.

Accepting Requests

The request managers and administrators can, at any point, accept *New* and *Handled* requests.

You cannot accept *Rejected* and *Finished* requests.

To accept a request

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. Tick the request(s) you want to accept.
3. In the menu, select **Edit > Accept**. The wizard **Status edition** opens.
4. In the field **Enter a note**, you can type in a reason for accepting or the user performing the task. This text is available on the request properties page, in the **Request history** panel.

5. Click on **[OK]** to complete the operation. The report opens and indicates the operation success.
6. Click on **[CLOSE]** to go back to the page Incoming requests.

Rejecting Requests

The request managers and administrators can at any point reject *New* and *Handled* requests.

You cannot reject *Accepted* and *Finished* requests.

To reject a request

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. Tick the request(s) you want to reject.
3. In the menu, select **✎ Edit > Reject**. The wizard **Status edition** opens.
4. In the field **Enter a note**, you can type in a reason for accepting or the user performing the task. This text is available on the request properties page **Request history** module.
5. Click on **[OK]** to complete the operation. The report opens and indicates the operation success.
6. Click on **[CLOSE]** to go back to the page Incoming requests.

Finishing Requests

Once the request has been dealt with, when the object has been added, edited or deleted, the request managers and administrators can set the requests to *Finished*.

Only *Accepted* requests can be finished.

To finish a request

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. Tick the request(s) you want to finish.
3. In the menu, select **✎ Edit > Finish**. The wizard **Status edition** opens.
4. In the field **Enter a note**, you can type in a reason for accepting or the user performing the task. This text is available on the request properties page **Request history** module.
5. Click on **[OK]** to complete the operation. The report opens and indicates the operation success.
6. Click on **[CLOSE]** to go back to the page Incoming requests.

Archiving Requests

Archiving a request actually means moving it to the Local Files Listing. This means that it is no longer listed on the Incoming requests and Outgoing request pages.

Archiving a request is useful for requests that have been dealt with, have been canceled or that were rejected. In any of these cases, once the requesting user has been informed, it is probably useless to keep the request in the list.

The request managers and administrators can archive *Canceled*, *Rejected* and *Finished* requests.

To archive a request

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.

2. Tick the request(s) you want to remove from the list.
3. In the menu, select  **Edit > Archive**. The wizard **Status edition** opens.
4. In the field **Local file name**, type in the name of the .txt file that contains the request details following the format *<file-name>.txt* .

By default, the file is saved in the directory */data1/exports* but you can save it in an existing sub-directory of */data1* if you want: specify the subdirectory name following the syntax *<sub-directory-name>/<file-name>.txt* .

5. Click on to complete the operation. The report opens and indicates the operation success.
6. Click on to go back to the page Incoming requests.

Administrating Requests Using Your Own Settings

You can customize the Workflow administration methods by editing some Workflow dedicated registry database entries. For more details, refer to the chapter [Customizing the Requests Administration](#).

Once you customized these entries, the restrictions detailed in the section [Administrating Requests Using the Default Statuses and Options](#) might not apply anymore. However, requests managers and administrators may still rely on the procedures detailed in said section to administer the requests from the page Incoming requests.

Chapter 71. Executing Requests

There are different ways of executing requests:

1. Use the *Execute* option from the page Incoming requests if you are using the Workflow default classes. For more details regarding this option, refer to the section [Executing Requests Using the Execute Option](#).
2. Use classes to integrate the requests to the addition, edition or deletion wizard. This method can be used if you use the default Workflow classes or if you use customized ones. For more details, refer to the section [Executing Requests Using Classes](#).
3. Go to the IPAM or DNS module and add, edit or delete the requested objects and change the status to Finished once the request was executed.

Executing Requests Using the Execute Option

If you are using the Workflow default classes listed in the chapter [Granting Access to Workflow Classes](#), you can use the option *Execute* to perform the action requested in the *New*, *Handled* and *Accepted* requests directly from the page Incoming requests.

To execute an addition request using the option Execute

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. In the column **Action**, filter the list using the key word *new* to only display requests for addition.
3. At the end of the line of the request for addition you want to execute, click on **Execute**. The wizard opens.
4. Depending on the classes configured you might have class dedicated pages. Select a class or none and click on **NEXT**.
5. On the object addition page, the object name and details are in a gray field as a reminder.

If the request demands a network addition, the wizard skips the network type selection page and directly creates the relevant network.
6. If need be, you can fill in the optional object details fields and configure advanced properties. Click on **NEXT**. The Workflow dedicated page opens.
7. In the drop-down list **Request**, the request you are executing is selected by default. The list can contain other request numbers if other requests for addition of a similar resource were created.
8. Under this field, the fields **Requested <object> name** and **Requestor motivation** contain the request original details as a reminder.
9. The requests for IP address addition have an extra page: the page **Aliases configuration**. You can add aliases if need be. Then click on **NEXT** to display the last page of the wizard.
10. Click on **OK** to complete the operation. The report opens and closes. The request status is now *Finished*, the object is now created.

To execute an edition request using the option Execute

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.

2. In the column **Action**, filter the list using the key word *modify* to only display requests for edition.
3. At the end of the line of the request for edition you want to execute, click on **Execute**. The wizard opens.
4. Depending on the classes configured you might have class dedicated pages. Select a class or none and click on **NEXT**.
5. On the object edition page, the object name and details are in a gray field as a reminder.
6. If need be, you can fill in the optional object details fields and configure default parameters. Click on **NEXT**. The Workflow dedicated page opens.
7. In the drop-down list **Request**, the request you are executing is selected by default. The list can contain other request numbers if other requests for edition of a similar resource were created.
8. Under this field, the fields **Requested <object> name** and **Requestor motivation** contain the request original details as a reminder.
9. The requests for IP address edition have an extra page: the page *Aliases configuration*. You can add aliases if need be. Then click on **NEXT** to display the last page of the wizard.
10. Click on **OK** to complete the operation. The report opens and closes. The request status is now *Finished*, the object is now edited.

To execute a deletion request using the option Execute

1. In the sidebar, go to **Workflow > Incoming requests**. The page **Incoming requests** opens.
2. In the column **Action**, filter the list to display only requests for deletion using the key word *delete*.
3. At the end of the line of the request for edition you want to execute, click on **Execute**. The wizard **Delete** opens.
4. The fields **Name**, **Address**, **Space name** and/or **DNS server name** contain the objects details as a reminder.
5. Click on **OK** to complete the operation. The report opens and closes. The request status is now *Finished*, the object is now deleted.

Once the request is executed, the requestor receives a notification email. The administrator or request manager can archive the request. For more details, refer to the section [Archiving Requests](#).

On the request properties pages, in the panel *Attached objects*, are listed all the object configuration details if the request concerned an addition or an edition. For instance, if a specific class or default parameters were set by the administrator or request manager.

Executing Requests Using Classes

You can use classes to automate the addition and edition requests execution. You cannot use them to automate the deletion requests. Using classes for the automation implies:

1. From Class Studio:
 - Creating a class applying to IPAM network, pool, address or DNS zone.
 - Adding the corresponding Pre-defined variable object to the class.
2. From the IPAM or DNS module:

- Applying the class when adding or editing the object.
 - At the end of the wizard, selecting the request matching the operation performed.
3. Archiving the request. For more details, refer to the section [Archiving Requests](#).

Configuring a Workflow Request Association Class

The classes that can automate the request execution apply to IPAM networks, pools, addresses or DNS zones.

If you do not already use a class for which you would like to add the Pre-defined variable, create a class. Otherwise, directly follow the procedure [To add a Workflow request association pre-defined variable](#).

To add a class to automate the Workflow request association

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the menu, click on **+ Add**. The wizard **Add a new class** opens.
4. In the field **Filename**, name your class. The name cannot contain any special characters. This field is mandatory.
5. In the field **Sub directory**, you can fill in the directory where you want to save your class. If it does not exist, it is created. On the wizards class selection page, classes placed in a directory are displayed as such: *<directory>/<class>*. This field is optional.
6. In the drop-down list **Module**, select *DNS* or *IPAM*.
7. In the drop-down list **Type**, select the resource of your choice: *DNS zone*, *Network*, *Pool* or *Address*.
8. In the section **Enable class**, tick the box.
9. Click on to complete the operation. The report opens and closes. The class is listed.

To add a Workflow request association pre-defined variable

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
4. In the class objects drop-down list, select *All Objects*. The class objects list refreshes.
5. In the list, click on . The wizard **Pre-defined variable** opens.
6. In the drop-down list **Name**, select the variable that suits your needs:
 - *WORKFLOW_ADD_TICKET_SPACE* to associate the class with space addition and/or edition requests.
 - *WORKFLOW_ADD_TICKET_BLOCK* to associate the class with non-terminal (block-type or subnet-type) network addition and/or edition requests.
 - *WORKFLOW_ADD_TICKET_SUBNET* to associate the class with terminal network addition and/or edition requests.
 - *WORKFLOW_ADD_TICKET_POOL* to associate the class with pool addition and/or edition requests.

- `WORKFLOW_ADD_TICKET_ADDRESS` to associate the class with address addition and/or edition requests.
 - `WORKFLOW_ADD_TICKET_DNSZONE` to associate the class with zone addition and/or edition requests.
7. In the field **Value**, type in the value `1` (one) to enable the variable.
 8. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

Once the class is configured, you can apply it from the DNS and/or IPAM module to automate the addition or edition of objects.

Applying a Workflow Request Association Class

To apply the Workflow request association class, it must be enabled and then selected in the addition or edition wizard.

To apply the Workflow request association class when adding an object

1. Depending on your needs, in the sidebar:
 - a. Go to **IPAM > Networks, Pools or Addresses**. The page opens.
 - b. Go to **DNS > Zones**. The page opens.
2. In the menu, click on **+ Add > <object>**. The corresponding wizard opens.
3. On the page **<Object> class**, select the class configured for the Workflow request association.
4. Configure the object according to your needs and click on **NEXT** until you get to the Workflow related page.
5. In the drop-down list **Ticket**, select an existing request for addition of the chosen object.
6. If you are adding an IP address, the page **Aliases configuration** opens. You can add aliases if need be. Then click on **NEXT** to display the last page of the wizard.
7. Click on **OK** to complete the operation. The report opens and closes. The object is listed. On the Workflow pages, the selected request is now *Finished*.

To apply the Workflow request association class when editing an object

1. Depending on your needs, in the sidebar:
 - a. Go to **IPAM > Networks, Pools or Addresses**. The page opens.
 - b. Go to **DNS > Zones**. The page opens.
2. Right-click over the object you want to edit. The contextual menu opens.
3. Click on **🔗**. The corresponding wizard opens.
4. On the **<page Object> class**, select the class configured for the Workflow request association.
5. Edit the object according to your needs and click on **NEXT** until you get to the Workflow related page.
6. In the drop-down list **Ticket**, select an existing request for addition of the chosen object.

7. If you are editing an IP address, the page **Aliases configuration** opens. You can add aliases if need be. Then click on **NEXT** to display the last page of the wizard.
8. Click on **OK** to complete the operation. The report opens and closes. The object is listed. On the Workflow pages, the selected request is now *Finished*.

Once the request is executed, the requestor receives a notification email. The administrator or request manager can archive the request. For more details, refer to the section [Archiving Requests](#).

Chapter 72. Customizing the Requests Administration

Depending on your needs, you can entirely customize the menu *Edit* of the page Incoming requests as well as the restrictions associated with the status edition. As detailed in the section [Administering Requests Using the Default Statuses and Options](#), you cannot set all the statuses to the requests as you please. As you can see in the image below.

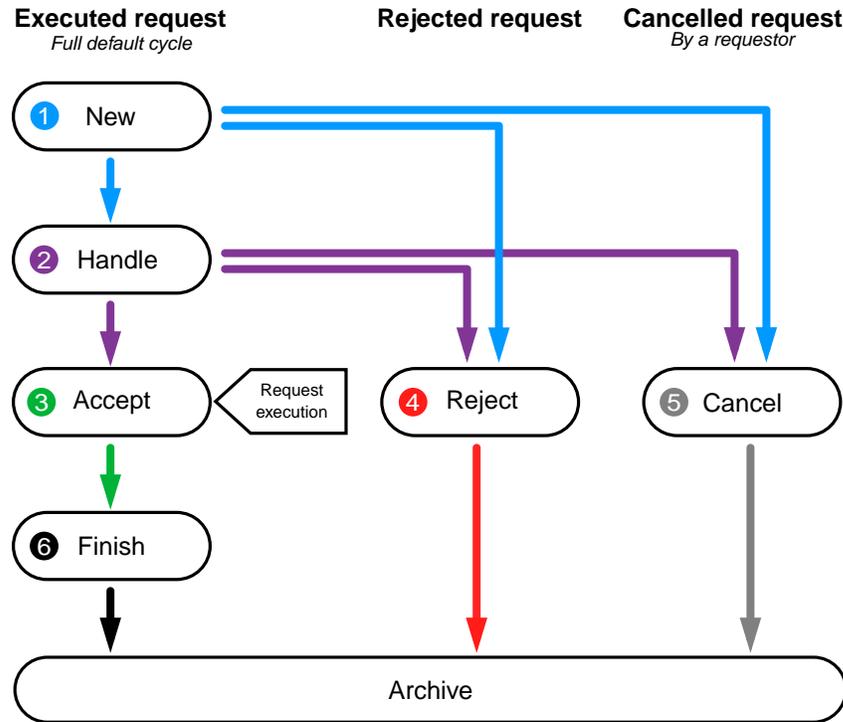


Figure 72.1. Workflow default status cycle

These default edition restrictions are all set in the registry database. The default configuration of the Workflow in the registry database is the following:

Table 72.1. Workflow registry database entries default value

Registry database entry	Default value
module.workflow.state.accept	incoming, wf-accept, t, nocallback, new-target,handle-operator
module.workflow.state.archive	incoming, wf-archive, t, archive_callback, cancel-admin, reject-admin,finish-admin, accept-operator
module.workflow.state.cancel	outgoing, wf-cancel, t, cancel_callback, new-source,handle-source
module.workflow.state.finish	incoming, wf-finish, t, finish_callback, accept-operator
module.workflow.state.handle	incoming, wf-handle, t, nocallback, new-target,handle-target
module.workflow.state.new	incoming, wf-new, f, nocallback, finish-admin
module.workflow.state.reject	incoming, wf-reject, t, nocallback, new-target,handle-operator
module.workflow.state_mail	new,handle,accept,reject,finish

You can edit default statuses, remove default statuses from the GUI and add new statuses. Whatever the customization you have in mind, we recommend that you take into consideration the section [Customized Statuses Best Practices](#).

Editing the Workflow Statuses

All the entries of the registry database dedicated to the Workflow configuration can be identified and filtered. You can edit these entries to suit internal processes, for instance skip some statuses that are obsolete to your organization or even grant more permissions to requestors and request managers.

Whether you decide to edit an existing status or hide it from the GUI, to make sure the request cycle is complete, we recommend that you follow, the sections [Status Edition Best Practices](#) and [Status Deletion Best Practices](#).

To edit the Workflow request statuses

1. Edit the status entries value:

Only users of the group *admin* can perform this operation.

- a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- c. In the search engine of the column **Name**, type in *module.workflow.state*. Only the Workflow status related keys are listed.
- d. In the column **Value**, click on the value of the entry you want to edit. The wizard **Registry database Edit a value** opens.
- e. In the field **Value**, specify the value of your choice following the [description of the Workflow Status Entries String](#) below.
- f. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

2. Register your changes:

- a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Expert**, click on **Register new macros & rules**. The wizard **Register all the latest macros and rules** opens.
- c. Click on **OK** to complete the operation. The report opens and works for a while. A notification pop-up appears in the lower right corner of the GUI when the operation is over.

The Workflow configuration entries are all named *module.workflow.state<detail>*. There are seven entries dedicated to the default statuses.

1. *module.workflow.state.accept* .
2. *module.workflow.state.archive* .
3. *module.workflow.state.cancel* .
4. *module.workflow.state.finish* .

5. `module.workflow.state.handle` .
6. `module.workflow.state.new` .
7. `module.workflow.state.reject` .

Each entry is important as it sets the permissions and restrictions related to the status. The status key value is a string in which the order matters. They must be separated by a coma as follows: `<page>`, `<icon>`, `<visibility>`, `<callback>`, `<attribute_1, attribute_2, ..., attribute_n>` .

As an example, the Accept status is detailed in the image below:

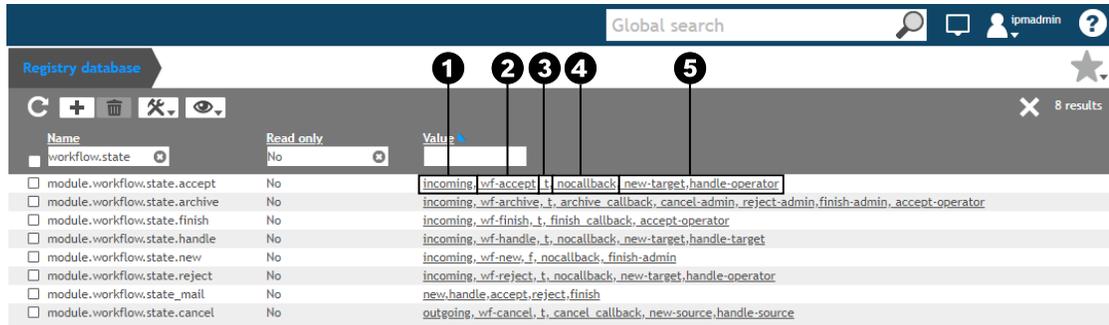


Figure 72.2. Structure of a workflow status registry entry

- ❶ Page where the status can be set.
- ❷ Icon preceding the *Accepted* status on the page.
- ❸ Visibility in the menu configuration.
- ❹ This section of the value is obsolete.
- ❺ Attributes defines who can set the Accept status and in what conditions.

In this example, the Accept status is displayed (*t*) on the page Incoming requests (*incoming*) and is preceded by the green icon. Any user with sufficient rights can accept *New* requests (*new-target*) and only the request manager who Handled or Rejected the request can accept it (*accept-operator*, *reject-operator*).

Each element of the string has a set of acceptable values that define the request status logic and organization that suits your needs:

Page

incoming specifies that the status is available on the page Incoming requests.

outgoing specifies that the status is available on the page Outgoing requests.

Icon

wf-accept allows to display the icon ❸, before the status name.

wf-archive does not display any icon as archiving means removing the request from the list.

wf-cancel allows to display the icon ❺, before the status name.

wf-finish allows to display the icon ❻, before the status name.

wf-handle allows to display the icon Ⓣ, before the status name.

wf-new allows to display the icon , before the status name.

wf-reject allows to display the icon , before the status name.

Visibility

t stands for true and indicates that the status is available in the menu *Edit* of the specified <page>.

f stands for false and indicates that the status is not displayed in the menu *Edit* of the specified <page>.

Callback parameters

This parameter is **obsolete**. You can find in the keys the values: *callback*, *nocallback*, *archive_callback* and *cancel_callback*. Do not edit them, they are part of the string.

Attributes

This last part of the string sets which user can set the status described in the string. This permission depends on who set the previous status: the user who set the status listed can now set the status described in the string.

The permissions structure follows the format: **<action>-<user>** in which action can be: **accept**, **archive**, **cancel**, **finish**, **handle**, **new** and **reject**, each one corresponds to the default statuses.

The users are:

- **admin** that is to any user in the group *admin*, including *ipmadmin*.
- **operator** the user that deals with the request. The other users belonging to the same group cannot perform the actions associated with *operator*: only the user who performed the action detailed in the status entry Value is the operator.
- **source** the user who created the request, i.e. the requestor.
- **target** any user with sufficient Workflow permissions, including *ipmadmin*. In other words, any user that can see the request in the list.

Therefore, only the users specified in the field Value of the status entry can set the status described and only if the previously set one of the statuses associated with their <user> name.

Editing the Email Notification Details

In addition to the status dedicated entries, there is one key dedicated to the notification of requestors via email: **module.workflow.state_mail**.

The default configuration sends an email to the requestors whenever the status request they created is edited. This is why, by default, it contains **new,handle,accept,reject,finish** .

The requestors only receive an email if their User profile was set properly. For more details, refer to the chapter [Managing Users](#) in the section [Adding Users](#) or [Editing Users](#).

To edit the email notification trigger details

1. Edit the status entries value:

Only users of the group *admin* can perform this operation.

- a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
 - c. In the search engine of the column **Name**, type in *module.workflow.state_mail*. Only this key is listed.
 - d. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
 - e. In the field **Value**, edit the content according to your needs. An email is sent to the requestor if the status attributed to a request they created is listed in the field. By default, the field contains *new,handle,accept,reject,finish*.
 - f. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.
2. Register your changes:
- a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **Expert**, click on **Register new macros & rules**. The wizard **Register all the latest macros and rules** opens.
 - c. Click on **OK** to complete the operation. The report opens and works for a while. A notification pop-up appears in the lower right corner of the GUI when the operation is over.

Adding a Workflow Status

You can add new statuses to the Incoming and Outgoing requests pages. This implies:

1. Adding the registry database entry following the Workflow entries format.
2. Translating the related menu option and status on the page.
3. Follow the [Status Addition Best Practices](#).

To add a Workflow request status

1. Add the status entry:

Only users of the group *admin* can perform this operation.

- a. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- c. Filter the list to only display the Workflow status related entries using the keyword: *module.workflow.state*.
- d. In the menu, click on **+ Add**. The wizard **Registry database Add an item** opens.
- e. In the field **Name**, type in the status name following the format: *module.workflow.state.<your-status-name>*.
- f. In the column **Value**, type in the characteristics of the new status following the format *<page>, <icon>, <visibility>, nocallback, <attribute_1>, <attribute_2>, <attribute_n>*. For more details, refer to the [description of the Workflow Status Entries String](#).

- g. Click on to complete the operation. The report opens and closes. The new entry is listed.
2. Register your changes:
 - a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **Expert**, click on **Register new macros & rules**. The wizard **Register all the latest macros and rules** opens.
 - c. Click on to complete the operation. The report opens and works for a while. A notification pop-up appears in the lower right corner of the GUI when the operation is over.

Once the entry is created and registered, the new status is visible in the menu *Edit* of the selected page as followed: *rq_<your-status-name>*. Once you attributed the status to a request, the request Status is *rq_in_<your-status-name>*. You can translate both using the page Language editor.

To translate the name of your Workflow statuses

1. From any page or wizard within SOLIDserver, copy the name of a field, page, column or menu that you want to replace with your label.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Language editor**. The page **Language editor** opens.
4. In the menu, click on **+ Add**. The wizard opens.
5. In the field **Key**, paste the status name. We recommend that you copy/paste the label name because Language editor is case sensitive.
6. If your appliance is displayed in English, in the field **English**, type in the new label value.
7. Click on to complete the operation. The entry is listed. Go back to the page where you copied the label to see the new name.

Customized Statuses Best Practices

Whether you decide to add statuses, edit statuses or remove them from the GUI, to complete the status customization we recommend that you follow the best practices below.

Status Addition Best Practices

Once you created a new status, you should:

- Edit the list **attributes** in the entries describing the statuses you do use. For instance, if you want to add a *Postpone* status that can be set after a request is accepted, you should add the *postpone-<user>* attribute in the value of the finish entry as well as the *accept-<user>* in the value of the postpone entry¹. For more details, refer to the [description of the Workflow Status Entries String](#) below.
- Edit all the status **icons** to make sure that the GUI respects your new request cycle.
- Add the status in the **email notification** entry. For more details, refer to the section [Editing the Email Notification Details](#).

¹This example is only valid if you still use the default statuses cycle.

- The new status cannot be executed using the *Execute* option as it only applies to New, Handled and Accepted requests. For more details, refer to the section [Executing Requests Using the Execute Option](#).
- The restrictions detailed in the [Administering Requests Using the Default Statuses and Options](#) no longer apply to your request status cycle.
- The request execution automation using the pre-defined variables class objects can still be configured. For more details, refer to the section [Executing Requests Using Classes](#).

Status Edition Best Practices

Once you decided to edit statuses or add new ones, keep in mind that:

- Once you edited the registry database entries the *Execute* option still only applies to New, Handled and Accepted requests. For more details, refer to the section [Executing Requests Using the Execute Option](#).
- The restrictions detailed in the [Administering Requests Using the Default Statuses and Options](#) no longer apply to your request status cycle.
- The request execution automation using the pre-defined variables class objects can still be configured. For more details, refer to the section [Executing Requests Using Classes](#).

Status Deletion Best Practices

To remove a status from the GUI, you recommend that you edit the status registry entry Value and set its visibility attribute to *f* (false). From then on the status is no longer visible in the menu *Edit*, and can no longer be used by users. Keep in mind that:

- Edit the list **attributes** in the entries describing the statuses you do use. For instance, if you want to remove the *Handle* status from the request management steps, you should remove all the *handle-<user>* attributes from the other statuses value field. For more details, refer to the [description of the Workflow Status Entries String](#) below.
- Edit all the status **icons** to make sure that the GUI respects your new request cycle.
- Remove the status from the **email notification** entry. For more details, refer to the section [Editing the Email Notification Details](#).
- Keep in mind that if the status was already set before you remove it from the menu, it is still displayed in the list.

Part XIII. Device Manager

Device Manager provides an overview of all the equipment on your network. Relying on both manual and automatic management options, it helps to piece together the data registered in the modules NetChange, IPAM, DHCP and DNS to map out the device interactions and their connections through interfaces and ports. It allows to organize devices and their content.

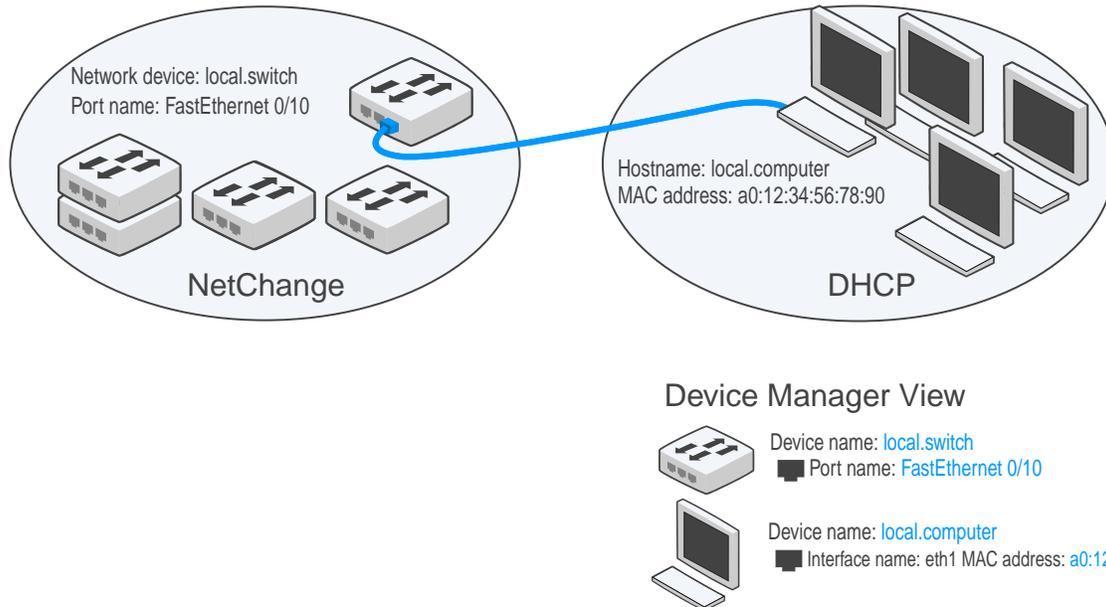


Figure 220. Information retrieved from NetChange and the DHCP

Note that all the data saved in Device Manager is never deleted, unless you decide to delete it. Therefore you can save a lot of information regarding users or pieces of equipment through their MAC address or IP address without impacting the other modules.

Device Manager hierarchy is composed of 2 levels:

- **Devices:** the highest level of the hierarchy. They contain ports and/or interfaces. For more details, refer to the chapter [Managing Devices](#).
- **Ports & interfaces:** the lowest level of the hierarchy. The ports are connected to other ports or interfaces, which allows to link the devices together on the network. For more details, refer to the chapter [Managing Ports and Interfaces](#).

In addition, the module allows to:

- **Replicate data to and from the IPAM.** Using the interfaces MAC address and/or IP address, you can automate updates. For more details, refer to the section [Managing the Interaction with the IPAM](#).
- **Automate updates from the DHCP.** A couple of rules allow to automatically update the database whenever relevant data is added in the DHCP. For more details, refer to the chapter [Rules Impacting Device Manager](#).

Note that from the module **Dashboards**, you can gather gadgets and charts on *Device Manager dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 73. Managing Devices

The devices are managed from the page **All devices**, where any device on the network can be managed (network devices, computers, virtual machines...) and uniquely identified based on the ports or interfaces it manages.

You can add them manually or automatically retrieve them. Devices can contain interfaces and/or ports, depending on the discovered MAC and IP addresses.

You can merge devices, duplicate them, edit their content or delete them. You cannot rename them.

Browsing Devices

Within the module Device Manager, the devices are the highest level of the hierarchy. It is required to create devices to manage ports and interfaces.

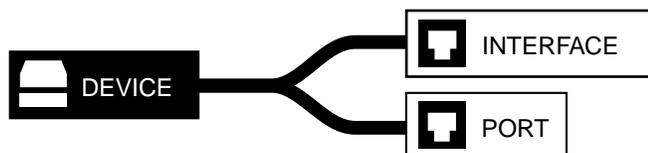


Figure 73.1. The device in Device Manager hierarchy

Browsing the Devices Database

To display the list of devices

1. In the sidebar, go to the **Device Manager** > **Devices**. The page **All devices** opens.
2. You can filter the list using the column search engines.

To display a device properties page

1. In the sidebar, go to the **Device Manager** > **Devices**. The page **All devices** opens.
2. At the end of the line of the device of your choice, click on . The properties page opens.

Customizing the Display on the Page All Devices

Users of the group *admin* can create customized column layouts. The button **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that a set of columns provides an overview of the devices interfaces and ports content:

- **Interfaces usage** and **Ports usage**: the total portion, in percent of used interfaces/ports on a device, along with a progression bar,
- **Number of Interfaces** and **Number of Ports**: the total number of interfaces/ports on the device,
- **Free Interfaces** and **Free Ports**: the number of available interfaces/ports on the device.

Note that the data listed in the column **IP Address** can be sorted but not filtered. It only retrieves and displays the IP address of all the interfaces of the device.

Managing the Devices Status and Visibility

Devices can be *Managed*, *Unmanaged* or *Imported*. Based on these statuses, you can filter the list from the column **Status** and, for example, only display the managed and imported devices using the value *!= Unmanaged* (different from Unmanaged).

The option *Manage* allows to manage or unmanage the device of your choice, whether you added it yourself or it was *Imported* from another module.

Keep in mind that you **cannot** unmanage a device associated with an IP address of the IPAM.

To manage/unmanage devices

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. Tick the device(s) of your choice.
3. In the menu, select  **Edit > Manage > Yes** or **No**. The wizard **Items management** opens.
4. Click on  to complete the operation. The report opens and closes. The device is marked  **Unmanaged** or  **Managed** in the column **Status**.

Adding Devices

You can add devices [manually](#) or [automatically](#) from several modules.

Before adding any device, we recommend [Configuring Device Manager](#) to make sure the data listed is consistent with the equipment configuration of your network.

Once added, you can decide which devices you want to display and deal with on the page *All devices*. For more details, refer to the section [Managing the Devices Visibility](#).

Keep in mind that Device Manager does not delete on its own the entries that you might delete in other modules. In that way, it provides an overview of former devices. To delete devices and their content refer to the section [Deleting Devices](#).

Note that you can also import devices from a CSV file from the page *All devices*. From then on, you can add or import the ports and interfaces it contains and organize your network as you please. For more details, refer to the section [Importing Data to Device Manager](#).

Configuring Device Manager

The option *Configure Device Manager* ensures the consistency of the links between the devices listed. You can execute it from both pages of the module.

We recommend configuring Device Manager before managing any device, port and interface because it compares the information of the other modules with what is listed in Device Manager. On the page *All ports & interfaces* of each device, the columns *Manually linked to* and *Automatically linked to* identify how the devices are linked together. If the column *Manually linked to* is empty, configuring Device Manager overwrites its content with the information collected during the automatic addition. This ensures that the data listed reflects the interaction between the devices on your network.

This option has to be ticked once. Afterward, a data check is performed each time a port or interface is added or edited.

To configure Device Manager automatic data check

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. In the menu, select  **Extra options > Configure Device Manager**. The wizard **Configure Device Manager reconciliation** opens.
3. Tick the box.
4. Click on  to complete the operation. The wizard closes, the page is visible again.

Automatically Adding Devices

A set of options automatically adds devices and the ports and interfaces they contain from Device Manager, NetChange and the IPAM.

Once you retrieved data automatically, you should monitor the column **Reconciliation** on the page *All ports & interfaces* to prevent any drift. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Note that you can also import devices from a CSV file from the page *All devices*. From then on, you can add or import the ports and interfaces it contains and organize your network as you please. For more details, refer to the section [Importing Data to Device Manager](#).

Automatically Adding Devices from the Page All Devices

After [Configuring Device Manager](#), you can execute the the option automatic discovery from the page *All devices*.

The option **Automatic discovery** performs a sweep of the other modules data and retrieves all the devices along with the ports and interfaces they contain. The option analyzes data in the modules NetChange, DHCP, IPAM and DNS. The more information there is in these modules, the more efficient is the option as it behaves as follows:

1. **It retrieves NetChange data:** network devices, ports and discovered items. The MAC addresses linked to the ports that have the interconnection set to *No* become interfaces in Device Manager. The discovered items DNS name is retrieved as well, the IP address is only retrieved if it is part of the IPAM. NetChange network devices can be managed as several devices in Device Manager depending on their content.

Note that if several ports in NetChange are linked to one MAC address, the option retrieves all the ports and the MAC address but randomly links it to one of the ports.

2. **It retrieves the names** of the devices, ports and interfaces if the information is available in the modules NetChange, IPAM, DHCP or DNS:
 - a. In **NetChange**, on the page *All discovered items*, the option retrieves all the MAC addresses, DNS names and IP addresses, only if they are part of the IPAM database.

If the MAC addresses retrieved have a DNS name, the option stops here.

If the MAC addresses do not have a DNS name, the option looks for it in the IPAM.

- b. In the **IPAM**, on the page *All addresses*, the option tries to match the collected MAC addresses with a name, using the available IPv4 and IPv6 addresses.

If a name is found for the MAC addresses, the option stops here.

If no name is found, the option looks for it in the DHCP.

- c. In the **DHCP**, on the page *All statics*, the option tries to match the collected MAC addresses with a name, using the available IPv4 statics and statics without IP.

If a name is found for the MAC addresses, the option stops here.

If no name is found, the option looks for it in the DNS.

- d. In the **DNS**, on the page *All RRs*, the option tries to match the collected IP addresses with a name, using the available A records.

If a name is found for the IP addresses, the option stops here.

If no name is found, Device Manager assigns the relevant MAC and IP addresses a name based on the information collected in all four modules.

3. It assigns a name to:

a. Devices

In Device Manager, the name of devices, including network devices, depends on their content.

- Network devices that contain one or several ports keep their NetChange name and only manage ports. The interfaces are managed in a different device.
- Devices which IP address is declared in an A record of the DNS are listed under that name in Device Manager.
- Network devices hosting the MAC address of one virtual machine, are named using the hostname of the virtual machine.
- Network devices hosting the MAC address of several virtual machines are named *vm_server_#*¹.
- Devices hosting an interface from which few information was retrieved are named *generic_#*.

b. Ports

All ports keep their NetChange name.

- Ethernet ports are named *Ethernet <slot>/<port>*, *FastEthernet <slot>/<port>*, *GigaEthernet <slot>/<port>* or *TenGigaEthernet <slot>/<port>*.
- Wifi ports are named *wifi#*, where # differentiates the ports belonging to one device.
- Virtual ports are named *Virtual port <slot>/<port>*.

c. Interfaces

All interfaces are named according to the port they are linked to. When a MAC address has a name, this name is used to name the device it belongs to.

- Interfaces linked to an ethernet port are named *eth#*, where # differentiates the interfaces belonging to one device.
- Interfaces linked to a wifi port are named *wifi#*.
- Interfaces linked to a virtual port are named *vm_interface_#*.

¹# is a number used to differentiate all these devices.

- Interfaces from which few information - except the MAC address - was retrieved are named *generic_#*.

To automatically add devices from the page **All devices**

1. In the sidebar, go to the  **Device Manager** > **Devices**. The page **All devices** opens.
2. In the menu, select  **Tools** > **Automatic Discovery**. The wizard **Automatic discovery** opens.
3. Click on to complete the operation. The report opens and closes. In the column **Status**, the devices are marked  **Imported** and you can manually manage and edit the devices, ports and interfaces.

Keep in mind that **the automatic discovery retrieves data but does not automatically update it**, if you delete data or make changes in NetChange, you must make the same changes in Device Manager.

If the automatic discovery created more devices that your need, you can merge devices to reorganize the data as you need. For more details, refer to the section [Merging Devices](#).

Automatically Adding Devices from NetChange

From the NetChange pages *All network devices* and *All discovered items* you can select objects to create devices, with the ports and interfaces they contain, in Device Manager.

Automatically Adding Network Devices in Device Manager

After [Configuring Device Manager](#), you can use the option **Create in Device Manager** on the page *All network devices* to select network devices and create them in Device Manager. It takes into account the information available on the pages *All network devices*, *All ports* and *All discovered items* and behaves as follows:

1. The device is created in Device Manager, it keeps the same name and contains all the network device ports, no matter their interconnection configuration.
2. The interfaces (discovered items) of the device are retrieved and based on the MAC address:
 - a. If the MAC address has a DNS name, a device is created using that name. The interface is named *eth#²*.
 - b. If the MAC address has no DNS name, a device is created, it is named *generic_#*. It contains the MAC address, it is named *eth#*.

If the option creates more devices that your need, you can merge devices to reorganize the data as you need. For more details, refer to the section [Merging Devices](#).

To automatically create Device Manager devices from NetChange

1. In the sidebar, go to  **NetChange** > **Network devices**. The page **All network devices** opens.
2. Tick the network device(s) you want to create in Device Manager.
3. In the menu, select  **Tools** > **Create in Device Manager**. The wizard **Create NetChange devices in Device Manager** opens.

²# differentiates interfaces within one device.

4. Click on **OK** to complete the operation. The report opens and closes and the page is visible again.
5. In the sidebar, go to **Device Manager > Devices**. The page **All devices** opens and the new devices are marked **Imported** in the column **Status**.

Automatically Adding Discovered Items in Device Manager

After [Configuring Device Manager](#), you can use the option **Populate Device Manager** on the page *All discovered items* to create interfaces from the discovered item of your choice. The MAC addresses selected create interfaces and devices as follows:

- If the selected MAC address has a DNS name, it creates an interface named *generic_#*³ that belongs to a device named using the DNS name.
- If the selected MAC address has no DNS name, it creates an interface named *generic_#* that belongs to a device named *generic_#*.

If the option creates more devices that your need, you can merge devices to reorganize the data as you need. For more details, refer to the section [Merging Devices](#).

To populate Device Manager with NetChange discovered items

1. In the sidebar, go to **NetChange > Discovered items**. The page **All discovered items** opens.
2. Tick the discovered item(s) you want to create in Device Manager.
3. In the menu, select **Tools > Populate Device Manager**. The wizard **Populate device manager** opens.
4. Click on **OK** to complete the operation. The report opens and closes and the page is visible again.
5. In the sidebar, go to **Device Manager > Devices**. The page **All devices** opens. The discovered item created a device that contains the interface. In the column **Status**, it is marked **Imported**.

Automatically Adding Devices from the IPAM

From the IPAM page *All addresses* you can automatically add devices managing interfaces created from the IPv4 and IPv6 addresses *In use* of your choice. The option names the objects as follows:

- The devices are named using the selected IP address name. If the IP addresses, v4 and/or v6, share the same name, like the addresses *Gateway*, they belong to one device. If the selected IP address does not have a name, the device is named *generic_#*.
- The interfaces are all named *generic_#*, where *#* differentiates interfaces within one device. Their IP address and the space they originated from are displayed in the column **Status**.

If the option creates more devices that your need, you can merge devices to reorganize the data as you need. For more details, refer to the section [Merging Devices](#).

To create devices and interfaces from the page All addresses

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.

³# differentiates interfaces within one device.

2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the *Used* IP addresses for which you want to create devices and interfaces.
4. In the menu, select **Tools > Populate Device Manager**. The wizard **Populate Device Manager** opens.
5. Click on **OK** to complete the operation. The report opens and works for a while and closes. The page is visible again. From the IPAM you can see the device and interface of the IP address in the columns *Device manager name* and *Device manager interface*.

In the module Device Manager, the device and interface are marked **Imported** in the column **Status**. In the column **IP address**, the IP address and the space they originated from are displayed as follows: *<ip-address> (<space>)*.

Manually Adding Devices

You can add a device manually in two different ways, either from the page *All devices* or when assigning an address from the IPAM.

Note that you can also import devices from a CSV file from the page *All devices*. From then on, you can add or import the ports and interfaces it contains and organize your network as you please. For more details, refer to the section [Importing Data to Device Manager](#).

Manually Adding Devices from the Page *All devices*

After [Configuring Device Manager](#), you can add a device manually from the page *All devices*.

To manually add a device

1. In the sidebar, go to the **Device Manager > Devices**. The page **All devices** opens.
2. In the menu, select **+ Add > Device**. The wizard **Add Device** opens.
3. If you or your administrator created classes at the device level, in the list **Device class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the field **Device**, name your device.
5. In the field **Description**, you can add a description.
6. Click on **OK** to complete the operation. The report opens and closes. The device is listed.

Note that, from the device addition wizard, you can also add the ports and/or interfaces it manages. For more details, refer to the section [Manually Adding Ports and Interfaces](#).

Manually Adding Devices from the IPAM

From the IPAM page *All addresses* you can assign an IPv4/IPv6 address and link it to a device and an interface directly from the IP address addition wizard. You can add as many devices and interfaces as you need.

First, you need to display the relevant advanced property via the wizard *Advanced properties customization*.

Table 73.1. IPAM / Device Manager advanced properties

In the wizard Advanced properties customization	In the addition/edition wizard
Enable to create devices from the IPAM	Create a device
	Device name
	Interface name

To display the device addition advanced properties at IP address level

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, select **Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
4. In the **IPAM / Device Manager interaction** section, tick the box **Enable to create devices from the IPAM**.
5. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Once you set the advanced properties display, the box *Create a device* allows to add new devices directly from the IP address addition and edition wizard both in IPv4 and IPv6.

To add a device when adding/editing an IP address

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit an IP address to display the Device Manager dedicated fields. The wizard opens. For more details regarding IP address addition or edition, refer to the chapter [Managing IP Addresses](#).
4. On the IP address configuration page of the wizard, configure the device addition advanced properties:
 - a. Tick the box **Create a device**, the fields *Device name* and *Interface name* appear. They are both required.
 - b. In the field **Device name**, type in the name of your new device.
 - c. In the field **Interface name**, type in the name of your new interface.
5. Click on **NEXT**. The page **Aliases configuration** opens.
6. Configure aliases if need be. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).
7. Click on **OK** to complete the operation. The report opens and closes.

The changes are visible on the IP address properties page, in the panel **Advanced properties**, and in the columns **Device manager name** and **Device manager interface**. Click on the device or interface name to access its properties page in Device Manager.

Once added from the IPAM module, you can edit the device and interface from Device Manager.

Duplicating Devices

At any time, you can duplicate the content and class parameters of a device. Duplicating devices can be used to anticipate provisioning.

For example, if you know that you will add a switch to NetChange that has a configuration similar to an existing device, you can duplicate the device and then edit it, etc. Once you the new switch is saved in Device Manager, you can complete the device content edition: you must update the MAC address of all your interfaces and the links between ports and devices. For more details, refer to the section [Editing Ports and Interfaces Properties](#).

If you know that you will add a new network device to NetChange and an existing device in Device Manager has configuration similar, you should duplicate the device to import more easily the network device to the page *All devices*. In this case you should:

1. Duplicate your device and name it like the coming NetChange network device, as detailed in the procedure below.
2. Edit the device content (number of ports and/or interfaces it manages, edit the links between the devices...) to make sure it matches the network device to come. For more details, refers to the sections [Adding Ports and Interfaces](#), [Editing Ports and Interfaces Properties](#) and [Deleting Ports and Interfaces](#).

To duplicate a device

1. In the sidebar, go to the  **Device Manager** > **Devices**. The page **All devices** opens.
2. Tick the device you want to duplicate. You can only duplicate one device at a time.
3. In the menu, select  **Edit** > **Duplicate**. The wizard **Duplicate device** opens.
4. In the field **Device name**, type in the name of the new device.
5. Click on to complete the operation. The report opens and closes. The device is listed. It contains the same ports and/or interfaces. However, the link from ports to device has to be set manually and interfaces MAC addresses are automatically generated.

Merging Devices

You can merge devices to manage the ports and interfaces they contain from a unique device.

Merging devices can be useful if you want to correct what was automatically found on the network. For instance, if after automatically retrieving information from NetChange, a port and an interface end up in two different devices even if they both belong to one laptop, you can merge these devices to manage them from a single device.

To merge devices

1. In the sidebar, go to the  **Device Manager** > **Devices**. The page **All devices** opens.
2. Tick the devices you want to merge.
3. In the menu, select  **Edit** > **Merge**. The wizard **Merge device** opens.
4. In the drop-down list **Name**, select the device that should include all the ports and interfaces. The other device(s) are emptied and deleted.

5. Click on **OK** to complete the operation. The report opens and closes. The device is listed, the other devices are no longer listed.

Deleting Devices

At any time you can delete a device, this also deletes the ports and interfaces it contains. This action is non-reversible.

To delete a device

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. Tick the device(s) you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The device and the ports and interfaces it contains are no longer listed.

Chapter 74. Managing Ports and Interfaces

From the page **All ports & interfaces**, you can manage the ports and the interfaces that belong to your devices.

You can add them to a specific device or when you create a device. You can also automatically retrieve them, along with the device they belong to, from other modules.

The ports link devices together, the interfaces are connected to the ports.

The interfaces have a MAC address and can have one or several IPv4 and/or IPv6 addresses. Both IPv4 and IPv6 addresses are listed on the page.

Note that, to minimize any error or distortion between what is really connected to the network and what is listed, you can track changes and reconcile data on this page. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Browsing Ports and Interfaces

Within Device Manager module, the ports and interfaces are the lowest level of the hierarchy. They both belong devices, and are listed on the same page.

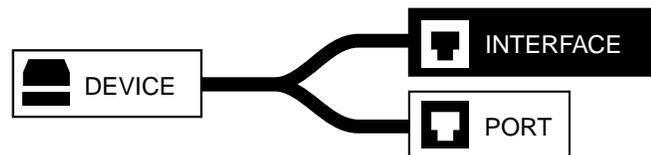


Figure 74.1. The interface in Device Manager hierarchy

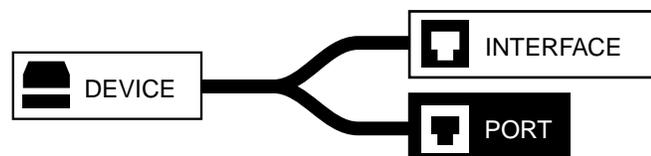


Figure 74.2. The port in Device Manager hierarchy

Browsing the Ports and Interfaces Database

To display the list of ports and interfaces

1. In the sidebar, go to **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
2. To display the ports & interfaces of a specific device:
 - a. In the column **Device name**, click on the name of the device of your choice. The main properties of the device are displayed.
 - b. In the breadcrumb, click on **All ports & interfaces**. The page **All ports & interfaces** opens.

To display a port or interface properties page

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
2. Filter the column **Type**, type in *port* or *interface* to list the objects that suit your needs.
3. At the end of the line of the port or interface of your choice, click on . The properties page opens.

Customizing the Display on the Page All Ports & Interfaces

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Keep in mind that the column **Addition date** provides extra information regarding the devices' content. You can use it to sort or filter the list.

Managing the Ports and Interfaces Status and Visibility

Ports and interfaces can be *Managed*, *Unmanaged* or *Imported*. Based on these statuses, you can filter the list from the column **Status** and, for example, only display the managed and imported ports and interfaces using the value *!= Unmanaged* (different from Unmanaged).

The option *Manage* allows to manage or unmanage the port or interface of your choice, whether you added it yourself or it was *Imported* from another module.

To manage/unmanage a device

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
2. Tick the port and/or interface of your choice. You can tick more than one.
3. In the menu, select  **Edit** > **Manage** > **Yes** or **No**. The wizard **Items management** opens.
4. Click on  to complete the operation. The report opens and closes. The device is marked  **Unmanaged** or  **Managed** in the column **Status**.

Adding Ports and Interfaces

You can add ports and interfaces [manually](#) or [automatically](#) retrieve them.

Before adding any port or interface, we recommend [Configuring Device Manager](#) to make sure the data listed is consistent with the equipment configuration of your network.

Once added, you can decide which items you display and deal with on the page, as detailed in the section [Managing the Ports and Interfaces Visibility](#).

Keep in mind that Device Manager does not delete the entries that you might delete in other modules on its own. In that way, it provides an overview of former port and interface interactions, to delete ports and interfaces refer to the section [Deleting Ports and Interfaces](#).

Note that you can also import ports and/or interfaces from a CSV file on the page *All ports & interfaces*. For more details, refer to the section [Importing Data to Device Manager](#).

Configuring Device Manager

The option *Configure Device Manager* ensures the consistency of the links between the devices on your network, through the ports and interfaces they manage.

This option has to be ticked once, **if you already configured Device Manager there is no need to do it again.**

On the page *All ports & interfaces* of each device, the columns *Manually linked to* and *Automatically linked to* identify how the devices are linked together. If the column *Manually linked to* is empty, configuring Device Manager overwrites its content with the information collected during the automatic addition. This ensures that the data listed reflects the interaction between the devices on your network.

To configure Device Manager automatic data check

1. In the sidebar, go to  **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
2. In the menu, select  **Extra options > Configure Device Manager**. The wizard **Configure Device Manager reconciliation** opens.
3. Tick the box.
4. Click on to complete the operation. The wizard closes, the page is visible again.

Automatically Adding Ports and Interfaces

You can automatically add ports and interfaces, along with the device they belong to when relevant, from Device Manager or the IPAM.

Once you retrieved data automatically, you should monitor the column **Reconciliation** on the page *All ports & interfaces* to prevent any drift. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Note that you can also import ports and/or interfaces from a CSV file on the page *All ports & interfaces*. For more details, refer to the section [Importing Data to Device Manager](#).

Automatically Adding Ports and Interfaces from the page All Devices

After [Configuring Device Manager](#) you can use the option **Automatic discovery** from the page *All devices* to retrieve devices, and the ports and interfaces they contain. The name of all the objects reflect the amount of information available in the modules NetChange, IPAM, DHCP and DNS.

This option creates devices from network devices. Some of these devices manage only ports and others manage one or several interfaces based on the MAC addresses retrieved.

On the page *All network devices*, the ports can have the following names:

- Ethernet ports can be named *Ethernet <slot>/<port>*, *FastEthernet <slot>/<port>*, *GigaEthernet <slot>/<port>* or *TenGigaEthernet <slot>/<port>*.
- Wifi ports are named *wifi#*, where # differentiates the ports belonging to one device.
- Virtual ports are named *Virtual port <slot>/<port>*.

On the page *All network devices*, the interfaces can have the following names:

- Interfaces linked to an ethernet port are named *eth#*, where # differentiates the interfaces belonging to one device
- Interfaces linked to a wifi port are named *wifi#*.
- Interfaces linked to a virtual port are named *vm_interface_#*.
- Interfaces from which few information - except the MAC address - was retrieved are named *generic_#*.

For more details regarding what data is retrieved and how it is named, refer to the section [Automatically Adding Devices from the Page All Devices](#).

To automatically add ports and interfaces from the page All devices

1. In the sidebar, go to the  **Device Manager** > **Devices**. The page **All devices** opens.
2. In the menu, select  **Tools** > **Automatic Discovery**. The wizard **Automatic discovery** opens.
3. Click on to complete the operation. The report opens and closes.
4. In the bread crumb, click on **All ports & interfaces**. The page opens.
5. In the column **Status**, the ports and interfaces are marked  **Imported**.

Once the automatic discovery retrieved ports and interfaces, you can rename or edit them if need be. For more details, refer to the section [Editing Ports and Interfaces Properties](#). You should also monitor the column **Reconciliation** as detailed in the section [Tracking Changes on the Page All ports & interfaces](#).

Keep in mind that the option retrieves data but does not automatically update it, if you delete data or make changes in NetChange, you must make the same changes in Device Manager.

Automatically Adding Interfaces from the IPAM

From the IPAM page *All addresses* you can automatically add interfaces and the devices they belong to from the IPv4 and IPv6 addresses *In use* of your choice. The option names the objects as follows:

- The devices are named using the selected IP address name. If the IP addresses, v4 and/or v6, share the same name, like the addresses *Gateway*, they belong to one device. If the selected IP address does not have a name, the device is named *generic_#*.
- The interfaces are all named *generic_#*, where # differentiates interfaces within one device. Their IP address and the space they originated from are displayed in the column **Status**.

No ports are added as the interfaces are created using the MAC address of the selected IP addresses.

To create devices and interfaces from the page All addresses

1. In the sidebar, go to  **IPAM** > **Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on or depending on your needs. The page refreshes and the button turns black.
3. Tick the *Used* IP addresses for which you want to create devices and interfaces.
4. In the menu, select  **Tools** > **Populate Device Manager**. The wizard **Populate Device Manager** opens.

- Click on **OK** to complete the operation. The report opens and works for a while and closes. The page is visible again. From the IPAM you can see the device and interface of the IP address in the columns *Device manager name* and *Device manager interface*.

In the module Device Manager, the device and interface are marked **Imported** in the column **Status**. In the column **IP address**, the IP address and the space they originated from are displayed as follows: *<ip-address> (<space>)*.

Manually Adding Ports and Interfaces

Manually adding ports and interfaces allows to correct what was found in NetChange or simply manage devices in accordance to your needs. You can add as many ports and interfaces as you want to a device to virtually link your devices.

Once you added data, you should monitor the column **Reconciliation** on the page *All ports & interfaces* to prevent any drift. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Note that you can also import ports and/or interfaces from a CSV file on the page *All ports & interfaces*. For more details, refer to the section [Importing Data to Device Manager](#).

Manually Adding Ports

After [Configuring Device Manager](#) you can add ports manually from the pages *All devices* and *All ports & interfaces*.

From the page *All devices* you can add as many ports as you want when you add a device. All these ports have the same name and can be numbered.

To manually add a device and the ports it contains

- In the sidebar, go to the **Device Manager > Devices**. The page **All devices** opens.
- In the menu, select **+ Add > Device**. The wizard **Add Device** opens.
- If you or your administrator created classes at the device level, in the list **Device class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Device**, name your device.
- Tick the box **Add port(s)/interface(s)**. The ports and interfaces configuration fields appear.
- In the drop-down list **Type**, select *Port*. The port related fields appear.
- Set the number of ports and their name:

Table 74.1. Port dedicated fields when adding a device

Field	Description
Name	Type in the port name. If you use the character #, it is replaced by a number no matter how many ports you want the add at once.
Number of ports	Type in the number of ports you want to add in the device you are creating. If you used # in their name, they are all numbered from 1 to <i>n</i> .

- You can link the port you are creating with another device port or interface. This is not required, if you do not want to link your port go to step 8.

Table 74.2. Options to link a port you are creating to a device

Field	Description
Link with device	Type in the name of the device you want to link the port with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.
Link with port/interface	Type in the name of the port or interface you want to link the port with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

- Click on **ADD**. The port is listed as such: *port: <number of ports> <port name>* in the **Interfaces/Ports** list. If you want to add more ports to the device. Repeat these actions for as many ports as needed.
- In the list **Interfaces/Ports**, you can set in which order the ports and interfaces are displayed selecting the items name and using the  and  buttons.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**. This automatically empties all the fields and allows to add other entries.

- Click on **OK** to complete the operation. The report opens and closes. The device is listed. On the pages *All devices* and *All ports & interfaces*, in the column **Status**, the device and the port(s) are marked  **Managed**.

From the page *All ports & interfaces*, you can add one port at a time in a specific device.

To manually add a port

- In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
- In the menu, click on **+ Add**. The wizard **Add port/interface** opens.
- In the drop-down list **Device**, select one of your existing devices.
- Click on **NEXT**. The next page opens.
- If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Name**, name the port.
- In the drop-down list **Type**, select *Port*.
- You can link the port you are creating with another device port or interface. This is not required, if you do not want to link your port go to step 8.

Table 74.3. Options to link a port you are creating to a device

Field	Description
Link with device	Type in the name of the device you want to link the port with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.

Field	Description
Link with port/interface	Type in the name of the port or interface you want to link the port with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

- Click on **OK** to complete the operation. The report opens and closes. The port is listed and marked **Managed**.

Manually Adding Interfaces

After [Configuring Device Manager](#) you can add interfaces from the pages *All devices* and *All ports & interfaces*.

Keep in mind that you can also add devices and the interfaces they contain from the IPAM. For more details, refer to the section [Manually Adding Devices from the IPAM](#).

From the page *All devices* you can add as many interfaces as you want when you add a device. Each interface has a unique name, they can have be assigned an existing IP address, they must have a MAC address that you either set or let Device Manager generate automatically, in this case the MAC address is not displayed on the page.

To manually add a device and the interfaces it contains

- In the sidebar, go to the **Device Manager > Devices**. The page **All devices** opens.
- In the menu, select **+ Add > Device**. The wizard **Add a Device** opens.
- If you or your administrator created classes at the device level, in the list **Device class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Device**, name your device.
- Tick the box **Add port(s)/interface(s)**. The ports and interfaces configuration fields appear.
- In the drop-down list **Type**, select *Interface*.
- In the field **Name**, type in the interface name.
- You can link the interface with an IP address. This step is optional.

Table 74.4. Options to link an interface you are creating with an IP address

Field	Description
MAC address	Type in the MAC address if you know it. You then have to type in the corresponding IP address.
IP Address	Type in a known IP address of the IPAM module, the corresponding MAC address is deduced and entered in the field MAC Address ^a .

^aIf the MAC address is already listed on the page *All ports & interfaces*, this interface addition is impossible.

- In the drop-down list **Space**, you can select one of the existing IPAM spaces.
- You can link the interface you are creating with another device port or interface. If you do not want to link the interface to a port, go to step 10.

Table 74.5. Options to link an interface you are creating to another device

Field	Description
Link with device	Type in the name of the device you want to link the interface with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.
Link with port/interface	Type in the name of the port you want to link the interface with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

- Click on **ADD**. In the list **Interfaces/Ports**, the interface is listed as such: *interface: <interface name> <MAC address> <IP Address>*. Repeat these actions for as many interfaces as you need.
- In the list **Interfaces/Ports**, you can set in which order the ports and interfaces are displayed selecting the items name and using the  and  buttons.

You can edit the entries of the list. Click on an entry, its configuration is displayed in the fields again, edit the entry and click on **UPDATE** or click on **DELETE** to remove an entry from the list. If you made changes that you do not want to save, click on **CANCEL**. This automatically empties all the fields and allows to add other entries.

- Click on **OK** to complete the operation. The report opens and closes. The device is listed. On the pages *All devices* and *All ports & interfaces*, in the column **Status**, the device and the interface(s) are marked  **Managed**.

From the page *All ports & interfaces*, you can add one interface at a time in a specific device.

To manually add an interface

- In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
- In the menu, click on **+ Add**. The wizard **Add port/interface** opens.
- In the drop-down list **Device**, select one of your existing devices. You can use your keyboard to find the device you are looking for.
- Click on **NEXT**. The next page, regarding ports and interfaces, opens.
- If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Name**, name the interface.
- In the drop-down list **Type**, select *Interface*. The interface related fields appear.
- You can link the interface with an IP address. This step is optional.

Table 74.6. Options to link an interface you are creating with an IP address

Field	Description
MAC address	Type in the MAC address if you know it. You then have to type in the corresponding IP address.

Field	Description
IP Address	Type in a known IP address of the IPAM module, the corresponding MAC address is deduced and entered in the field MAC Address ^a .

^aIf the MAC address is already listed on the page All ports & interfaces, this interface addition is impossible.

- In the drop-down list **Space**, you can select one of the existing IPAM spaces.
- You can link the interface you are creating with another device port or interface. If you do not want to link the interface to a port, go to step 10.

Table 74.7. Options to link an interface you are creating to another device

Field	Description
Link with device	Type in the name of the device you want to link the interface with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.
Link with port/interface	Type in the name of the port you want to link the interface with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

- Click on to complete the operation. The report opens and closes. The interface is listed and marked  **Managed**.

Editing Ports and Interfaces Properties

To edit the ports and interfaces implies:

- Renaming a port or an interface.
- Editing the links between existing ports/interfaces and devices.
- Editing the links between existing interfaces and devices, and the interface MAC address.

Editing port and interfaces is required if you duplicated existing devices. Once you duplicated the relevant device, you must edit its content as follows:

- Add/delete the ports and interfaces it contains as needed. For more details, refer to the sections [Adding Ports and Interfaces](#) and [Deleting Ports and Interfaces](#).
- Manually link the ports/interfaces to the needed device as detailed in the sections [Editing a Port](#) and [Editing an Interface](#).

Renaming a Port or Interface

You can change the name of a port or an interface name from its properties page or from the contextual menu on the page *All port & interfaces*.

Keep in mind that ports retrieved from NetChange had a name before you chose to manage them in Device Manager. Once you renamed a port, both NetChange and Device Manager names are displayed on its properties page: the field *Name* displays your name, the field *NetChange port name* displays the port original name.

To rename a port from the page All ports & interfaces

- In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.

2. In the **Type** column, type in *port* to only display the ports.
3. Right-click over the **Name** of the port you want to edit. The contextual menu opens.
4. Click on . The wizard **Edit a port or interface** opens.
5. If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Name**, rename the port.
7. Click on **OK** to complete the operation. The report opens and closes. Your modified port name is listed, its former name is no longer visible.

To rename an interface from its properties page

1. In the sidebar, go to  **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
2. In the **Type** column, type in *interface* to only display the interfaces.
3. Filter the list if needed.
4. Click on the name of the interface you want to edit. The properties page opens.
5. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a port or interface** opens.
6. If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **Name**, rename the interface.
8. Click on **OK** to complete the operation. The report opens and closes. In the **Main properties** panel, the name is edited.

Editing a Port

At any time, the ports' links toward devices can be edited.

If you duplicated a device, you must edit these links. Once a device is duplicated, the newly created device ports are not linked to any other device. In this case, you have to create the link between the ports and the needed device port or interface. To successfully edit the port links between devices, you must:

1. Add a new link toward the newly created device interfaces.
2. Perform an automatic discovery if both the ports and interfaces links are edited.
3. Check the data.

To link a port to another device's interface

1. In the sidebar, go to  **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.

2. Filter the names to display the port name if needed.
3. Click on the name of the port you want to link to another device interface. The port properties page opens.
4. In the panel **Main properties**, click on **[EDIT]**. The wizard **Edit a port or interface** opens.
5. If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **[NEXT]**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Specify the other device port or interface.

Table 74.8. Options to link a port/interface you are editing to another device

Field	Description
Link with device	Type in the name of the device you want to link the port with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.
Link with port/interface	Type in the name of the port or interface you want to link the port with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

7. Click on **[OK]** to complete the operation. The report opens and closes. The device you selected is visible in the panel **Main properties** in the *Manually linked to* line, the selected interface is between brackets. If you go back to the page **All ports & interfaces**, you have the same information in the column *Manually linked to*.

Now that the links are saved, if you already added the new device in NetChange, you can run the automatic discovery. For more details, refer to the section [Adding Network Devices](#).

If you also have interfaces in that device, edit their links and MAC addresses before running the automatic discovery. For more details, refer to the section [Editing an Interface](#).

To automatically add ports and interfaces

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. In the menu, select  **Tools > Automatic Discovery**. The wizard **Automatic discovery** opens.
3. Click on **[OK]** to complete the operation. The report opens and closes. The devices are all listed, their content and names depend on what was found. For more details, refer to the section [Automatically Adding Devices from the Page All Devices](#).

Once your changes are done and the list of ports is up-to-date, you can compare the data added manually and automatically.

To check the automatic discovery results

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. Click on the **Name** of the new device. The page **All ports & interfaces** of the device opens.
3. Compare the content columns **Manually linked to** and **Automatically linked to**.
4. Make sure there is no *drift* in the column **Reconciliation**. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Editing an Interface

At any time, the interfaces' links toward devices can be edited.

If you duplicated a device, you must edit the links. Once a device is duplicated, the newly created device interfaces are not linked to any other device and their MAC address is probably incorrect. In this case, you have to create the link between the interfaces and the needed device port or interface. To successfully edit the interfaces' links between devices, you must:

1. Add a new link toward the newly created device interfaces.
2. Update the MAC address of the interface.
3. Perform an automatic discovery if both the ports and interfaces links are edited.
4. Check the data.

To link an interface to another device's interface

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
2. Filter the names to display the interface name if needed.
3. Click on the name of the interface you want to link to another device interface. The port properties page opens.
4. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a port or interface** opens.
5. If you or your administrator created classes at the port and interface level, in the list **Host interface class** select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Specify the other device port or interface.

Table 74.9. Options to link an interface to another device's interface

Field	Description
Link with device	Type in the name of the device you want to link the interface with. The auto-completion retrieves a list of existing devices matching this name that you can choose from.
Link with port/interface	Type in the name of the interface you want to link the interface with. The auto-completion retrieves a list of available ports and interfaces matching this name that you can choose from.

7. Click on **OK** to complete the operation. The report opens and closes. The device you selected is visible in the panel **Main properties** in the *Manually linked to* line, the selected interface is between brackets. If you go back to the page **All ports & interfaces**, you have the same information in the *Manually linked to* column.

Once you linked the interfaces to another device interface, update its MAC address.

To update a MAC address

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.

2. Tick the interface which MAC address you want to update. Filter the data if needed.
3. In the menu, select  **Edit > Update MAC**. The wizard **Update mac address** opens.
4. In the field **MAC address**, type in the new MAC address.
5. Click on to complete the operation. The report opens and closes. The interface is listed with the new MAC address. The MAC address is also updated within the IPAM module.

Once you updated a MAC address, the former MAC address is deleted and the IP address(es) it is linked to are saved whether it is an IPv4 or an IPv6 address.

Now that the links are saved and the MAC addresses updated, if you already added the new device in NetChange, you can run the automatic discovery. For more details, refer to the section [Adding Network Devices](#).

If you also have ports in that device, edit their links as well before running the automatic discovery. For more details, refer to the section [Editing a Port](#).

To automatically add ports and interfaces

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. In the menu, select  **Tools > Automatic Discovery**. The wizard **Automatic discovery** opens.
3. Click on to complete the operation. The report opens and closes. The devices are all listed, their content and names depend on what was found. For more details, refer to the section [Automatically Adding Devices from the Page All Devices](#).

Once your changes are done and the list of interfaces is up-to-date, you can compare the data added manually and automatically.

To check the automatic discovery results

1. In the sidebar, go to the  **Device Manager > Devices**. The page **All devices** opens.
2. Click on the **Name** of the new device. The page **All ports & interfaces** of the device opens.
3. Compare the content columns **Manually linked to** and **Automatically linked to**.
4. Make sure there is no *drift* in the column **Reconciliation**. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Tracking Changes on the Page All ports & interfaces

To minimize the risk of saving inaccurate information in the module, the page *All ports & interfaces* allows to track changes via a column and an option, both called **reconciliation**.

The reconciliation allows to compare the manual and/or automatic links between devices through ports and interfaces.

The Column Reconciliation

The column *Reconciliation* compares the data entered in the columns *Automatically linked to* and *Manually linked to*. Note that:

- It works in close relation with the reconciliation option. For more details, refer to the section [The Reconciliation Option](#).

- The data retrieved automatically always has the upper hand in Device Manager so do not use the reconciliation option if you know that what you entered manually corresponds to the way you want to manage your items.
- Editing the devices topology from the IPAM changes the content of the column *Manually linked to*. For more details, refer to the section [Editing the Devices Topology from the IPAM](#).

The column can contain the following:

Table 74.10. Possible content of the column Reconciliation

Type	Description
✔ OK	The information displayed in the columns <i>Automatically linked to</i> and <i>Manually linked to</i> is a match.
⊗ N/A	No information is displayed in either column.
▲ Drift	The information displayed in the columns <i>Automatically linked to</i> and <i>Manually linked to</i> is not a match.

The Top List *Alert on ports/interfaces reconciliation drift* tracks any drift in the column. For more details, refer to the section [Gadgets Displayed by Default](#).

The Reconciliation Option

The reconciliation option is here to proofread the link created manually versus the data entered automatically. The first way to avoid getting any *Drift* is to configure Device Manager before adding any items to it. Indeed, with this option both *Automatically linked to* and *Manually linked to* should basically contain the exact same data.

However, if you decided to enter some data manually, you can reconcile both link related columns with the reconciliation option. That is to say, the content of the column *Manually linked to* overwrites the content of the column *Automatically linked to*.

To use the reconciliation option to save the manual device links

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
2. Filter data if needed. For instance, through **Drift** in the **Reconciliation** column.
3. Tick the port(s) and/or interface(s) you want to reconcile.
4. In the menu, select  **Edit** > **Reconcile**. The wizard **Reconciliation** opens.
5. Click on  to complete the operation. The report opens and closes. The items disappear from the list if the **Reconciliation** column was filtered by *Drift* as the value of the selected items is now *OK*.

Deleting Ports and Interfaces

You can delete interfaces and ports from the page All ports & interfaces. Keep in mind that:

- **The ports and interfaces deletion is manual and non-reversible.**
- **Deleting interfaces linked with the IPAM, breaks the link between the IP addresses and the device.**

To delete ports or interfaces

1. In the sidebar, go to  **Device Manager** > **Ports & interfaces**. The page **All ports & interfaces** opens.
2. Tick the port(s) and/or interface(s) you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on  to complete the operation. The report opens and closes. The selected item(s) is/are no longer listed.

Chapter 75. Managing the Interaction with the IPAM

Device Manager offers the possibility to display IPv4 and IPv6 addresses on the same list. Like the IPAM and the DHCP modules, it uses the Dual Stack protocol. Which is why you have the possibility to assign IP addresses in both versions of the Internet Protocol to one interface.

The automatic discovery option provides an automated assignment and display of both IPv4 and IPv6 addresses. It requires the IP address and interface MAC address to be associated in the IPAM module prior to running the option. Note that the MAC address in question should be part of NetChange discovered items list. Once you made changes in either module, you can run the option again following the procedure [To automatically add devices from the page All devices](#).

From the IPAM module, a set of options allows to edit Device Manager database. You can:

- **Add devices from the page All addresses** when assigning an IP address. For more details, refer to the section [Manually Adding Devices from the IPAM](#).
- **Associate IP addresses to existing interface or remove that link**. For more details, refer to the section [Managing the IP Addresses / Interfaces Link from the IPAM](#).
- **Edit the link between devices from the page All addresses**. For more details, refer to the section [Editing the Devices Topology from the IPAM](#).

From Device Manager, setting the MAC address of an interface also updates the IPAM database. For more details, refer to the section [Assigning IP Addresses to an Interface Using their MAC Address](#) below.

Assigning IP Addresses to an Interface Using their MAC Address

At any time, you can manually assign IPv4 and IPv6 addresses to existing interfaces provided that any relevant piece of information is already saved in SOLIDserver database (in the IPAM and NetChange). For that reason, adding, removing or editing an IP address MAC address, might change or remove an existing IP address/interface link.

Assigning IPv4 Addresses to an Interface

The IP addresses of newly added IPAM networks are not taken into account by Device Manager unless you assign them. Therefore, you have to manually assign the MAC addresses to the IP addresses.

To assign an IPv4 address to an interface using its MAC address

1. **Within Device Manager**
 - a. In the sidebar, go to  **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
 - b. Order the list by MAC address. The interfaces are listed first.
 - c. Right-click over the **Name** of the interface of your choice. The contextual menu opens.

- d. Click on . The interface properties page opens.
- e. In the panel **Main properties**, copy the MAC address.

2. Within the IPAM module

- a. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
- b. Click on an available address. The pop-up message *This address is free, do you want to assign it?* opens.
- c. Click on **[OK]**. The wizard opens.
- d. If you or your administrator created classes, the list **IP address class** is visible. Select a class or *None* and click on **[NEXT]**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- e. The field **IP address name** is gray and empty.
- f. The field **IP address** displays the IP address.
- g. In the field **MAC address**, paste your MAC address.
- h. In the field **Shortname**, name your IP address: it is automatically displayed in the field **IP address name**.
- i. Click on **[NEXT]**. The page **Aliases configuration** opens. There is nothing to set up in this page when you are simply assigning an IP address to an interface. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).
- j. Click on **[OK]** to complete the operation. The page IPv4 addresses is visible again and the IP address is listed as *Used*, named and has a MAC address.

If the report page displays the warning message *MAC address already used. (Space: ..., Address:)*, on as many lines as IP address(es) used on the interface, click on **[OK]** to commit the addition of the extra IP address on the interface. To cancel the assignment, click on **[CLOSE]**. To edit the MAC address, click on **[PREVIOUS]**. The *Aliases configuration* page opens first, click on **[PREVIOUS]** again to open the page *Add an IPv4 address* where you can make the needed changes.

Assigning IPv6 Addresses to an Interface

With IPv6 addresses, the interface assignment also involves going through the IPAM module and editing the interface within Device Manager. However, we recommend that you make sure the MAC address of the interface using this IPv6 address is present among the NetChange discovered items. That way it is directly detected by Device Manager.

To assign an IPv6 address to an interface using its MAC address

1. Within Device Manager

- a. In the sidebar, go to **Device Manager > Ports & interfaces**. The page **All ports & interfaces** opens.
- b. Order the list by MAC address. The interfaces are listed first.
- c. Right-click over the **Name** of the interface of your choice. The contextual menu opens.

- d. Click on . The interface properties page opens.
- e. In the panel **Main properties**, copy the MAC address.

2. Within NetChange

- a. In the sidebar, go to  **NetChange > Discovered items**. The page **All discovered items** opens.
- b. In the search engine of the column **Interco**, click on  to remove the default filter.
- c. In the search engine of the column **MAC Address**, paste your address to make sure it is part of NetChange items.
- d. On your keyboard, hit Enter. If the MAC address is listed, go to the next step. If it is not listed, go back to *step 1* and find an interface that is part of the list All discovered items.

3. Within the IPAM module

- a. In the sidebar, go to  **IPAM > Addresses**. The page **All addresses** opens.
- b. On the right-end side of the menu, click on . The page refreshes and the button turns black.
- c. Click on an available address. The pop-up message *This address is free, do you want to assign it?* opens.
- d. Click on . The wizard **Add an IPv6 address** opens.
- e. If you or your administrator created classes, the list **IP address class** is visible. Select a class or *None* and click on . The next page of the wizard opens.
- f. The field **IP address name** is gray and empty.
- g. The field **IP address** displays the IP address.
- h. In the field **MAC address**, paste your MAC address.
- i. In the field **Shortname**, name your IP address: it is automatically displayed in the **IP address name** field.
- j. Click on . The page **Aliases configuration** opens. There is nothing to set up in this page when you are simply assigning an IPv6 address to an interface. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).
- k. Click on  to complete the operation. The report opens and closes. The IPv6 addresses list opens again and the IP address is listed as used, named and has a MAC address.

If the report page displays the warning message *MAC address already used. (Space: ..., Address: ...)*, on as many lines as IP address(es) used on the interface, click on  to commit the addition of the extra IP address on the interface. To cancel the assignment, click on . To modify the MAC address, click on . The *Aliases configuration* page opens first, click on  again to open the page *Add an IPv6 address* where you can make the needed changes.

4. Within Device Manager

- a. In the sidebar, go back to  **Device Manager**. The interface properties page opens.
- b. Click on . The wizard **Edit a port/interface** opens.

- c. Click on **OK** to complete the operation. The report opens and closes. In the panel **Interface attachments**, the IPAM section regarding v6 addresses is updated and display the new IP address information. The address is visible in the list **All ports & interfaces**.

You can use the procedure above for as many IP addresses as needed for one interface. Beyond one IPv6 address, the addition wizard displays a report step listing the IP addresses already used on this interface to make sure that you actually want to use an extra IP address.

Managing the IP Addresses / Interfaces Link from the IPAM

From the page All addresses, you can associate IP addresses to existing interfaces. You can either use advanced properties to provide a link between them or a dedicated option in the menu *Edit*. This menu also provides an easy way to break the link between an IP address and an interface.

The columns **Device manager name** and **Device manager interface** allow to display the interactions on that page. For more details regarding columns display, refer to the section [Customizing the List Layout](#).

Any ports and interfaces changes made from the IPAM change the content of the column *Manually linked to*. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Using Advanced Properties to Associate IP Addresses with Interfaces

You can associate IP addresses to existing interfaces directly from the IPAM page *All addresses*.

First, you need to display the relevant advanced property via the wizard *Advanced properties customization*.

Table 75.1. IPAM / Device Manager advanced properties

In the wizard Advanced properties customization	In the addition/edition wizard
Enable to link IP addresses with existing devices	Device name
	Interface name

To display the IP address/interface link advanced properties at IP address level

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, select **Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
4. In the **IPAM / Device Manager interaction** section, tick the box **Enable to link IP addresses with existing devices**.
5. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Once you set the advanced properties display, the fields *Device name* and *Interface name* allow to link the IPv4 or IPv6 addresses you add or edit with existing interfaces.

To link an IP address with an interface using advanced properties

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit an IP address to display the Device Manager dedicated fields. The wizard opens. For more details regarding IP address addition or edition, refer to the chapter [Managing IP Addresses](#).
4. On the IP address configuration page of the wizard, configure the device addition advanced properties:
 - a. In the field **Device name**, type in the name or part of the name of an existing device. The auto-completion retrieves a list of devices matching this name that you can choose from.
 - b. In the field **Interface name**, type in the name or part of the name of an existing interface. The auto-completion retrieves a list of interfaces matching this name that you can choose from. Once you selected an interface, its name is displayed as follows: *<interface name> (<device name> - <number of IP addresses associated with the interface>)*.

If you do not specify an interface, the IP address is only associated with the device, whose name is displayed on the IP address properties page.
5. Click on **NEXT**. The page **Aliases configuration** opens.
6. Configure aliases if need be. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).
7. Click on **OK** to complete the operation. The report opens and closes.

The changes are visible on the IP address properties page, in the panel **Advanced properties**, and in the columns **Device manager name** and **Device manager interface**. Click on the device or interface name to access its properties page in Device Manager.

Using the Menu to Manage the IP Addresses / Interfaces Link

From the page All addresses menu, you can set up, edit or remove links between your IP addresses and an existing Device Manager interface.

Linking IP Addresses with Existing Interfaces

Once you assigned IP addresses, you can link them to the existing device and interface of your choice through the menu. The wizard also allows you to edit, i.e. overwrite, an existing link between the IP address and an interface.

The auto-completion provided in the device name and interface name only lists the device and interfaces marked as *Managed* and *Imported*. The *Unmanaged* items are not listed.

To link an IP address and an interface using the menu

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the IP address(es) of your choice.

4. In the menu, select  **Edit > Link IP addresses to Device Manager interfaces**. The wizard **Link IP addresses to Device Manager interfaces** opens.
5. Set the link following the table below.

Table 75.2. IP Address / Interface link configuration

Field	Description
Device name	Type in the name or part of the name of an existing device. The auto-completion retrieves a list of devices matching this name that you can choose from.
Interface name	Type in the name or part of the name of an existing interface. The auto-completion retrieves a list of interfaces matching this name that you can choose from. Once you selected an interface, its name is displayed as follows: <i><interface name> (<device name> - <number of IP addresses associated with the interface>)</i> .
Overwrite	Tick this box to edit an existing link and overwrite it with a link to the device and interface you specified in the above fields.

6. Click on  to complete the operation. The report opens and closes. The list **All addresses** is visible again. The changes are visible in the dedicated columns, on the IP address properties page and in Device Manager.

Removing the Link Between IP Addresses and Interfaces

Once you set a link between an IP address and an interface, you can remove it using the menu **provided that no MAC address was used when assigning the IP address**. If your IP address was assigned a MAC address, you need to edit the IP address, remove the MAC address and then follow the procedure below to remove the link with the interface. For more details regarding IP address edition, refer to the section [Editing an IP Address](#).

To remove an IP address / interface link using the menu

1. In the sidebar, go to  **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. Tick the IP address(es) of your choice.
4. In the menu, select  **Edit > Remove IP addresses / Device Manager interfaces link**. The wizard **Link IP addresses to Device Manager interfaces** opens.
5. Set the link following the table below.

Table 75.3. IP Address / Interface link configuration

Field	Description
Device name	Type in the name or part of the name of an existing device. The auto-completion retrieves a list of devices matching this name that you can choose from.
Interface name	Type in the name or part of the name of an existing interface. The auto-completion retrieves a list of interfaces matching this name that you can choose from. Once you selected an interface, its name is displayed as follows: <i><interface name> (<device name> - <number of IP addresses associated with the interface>)</i> .
Overwrite	Tick this box to edit an existing link and overwrite it with a link to the device and interface you specified in the above fields.

6. Click on  to complete the operation. The report opens and closes. The list **All addresses** is visible again. The changes are visible in the dedicated columns, on the IP address properties page and in Device Manager.

Editing the Link Between IP Addresses and Interfaces

Once you set up a link between an IP address and an interface, you can edit it to link the IP address with a different interface.

First, as detailed in the section [Linking IP Addresses with Existing Interfaces](#), you can simply specify a device and interface and tick the box **Overwrite**.

Second, if the IP address was assigned a MAC address, you can simply edit the MAC address to link the IP address with another interface. For more details regarding IP address edition, refer to the section [Editing an IP Address](#).

Editing the Devices Topology from the IPAM

You can manage the devices topology from the IPAM page All addresses, that is to say edit the link between an interface with another device port.

Note that:

- When adding an IP address, you need to tick the device addition or association property to allow users to first set the link between the IP address and the interface and then see the *Link with device* and *Link with interface* fields.
- When editing an IP address, the related fields are displayed when you edit an IP address **only if the IP address is already associated with a device and an interface**.

First, you need to enable the relevant advanced property via the wizard *Advanced properties customization*.

Table 75.4. IPAM / Device Manager advanced properties

In the wizard Advanced properties customization	In the addition/edition wizard
Enable to edit the devices topology from the IPAM	Link with device
	Link with port

To display the devices link advanced properties at IP address level

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, select **Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
4. In the **IPAM / Device Manager interaction** section, tick the box **Enable to edit the devices topology from the IPAM**.
5. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

Once you set the advanced properties display, the fields *Link with device* and *Link with port* allow to link the IPv4 or IPv6 addresses you are adding or editing with any port on the device of your choice.

To link devices from the IPAM

1. In the sidebar, go to **IPAM > Addresses**. The page **All addresses** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. To set a link between devices, add an IP address. The wizard opens. For more details regarding IP address addition or edition, refer to the chapter [Managing IP Addresses](#).
4. To edit the topology, edit an IP address already associated with an interface to display the Device Manager dedicated fields. The wizard opens. For more details regarding IP address addition or edition, refer to the chapter [Managing IP Addresses](#).
5. On the IP address configuration page of the wizard, configure the device addition advanced properties:
 - a. In the field **Link with device**, type in the name of an existing device. The auto-completion retrieves a list of devices matching this name.
 - b. In the field **Link with port**, type in the name of an existing port. The auto-completion retrieves a list of ports matching this name. Once you selected an interface, its name is displayed in the field along with the device.
6. Click on **NEXT**. The page **Aliases configuration** opens.
7. Configure aliases if need be. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).
8. Click on **OK** to complete the operation. The report opens and closes.

The changes are visible on the page **All ports & interfaces**, in the column **Manually linked to**. Therefore, you might need to reconcile the content of the columns *Automatically linked to* and *Manually linked to*. For more details, refer to the section [Tracking Changes on the Page All ports & interfaces](#).

Chapter 76. Rules Impacting Device Manager

From the module Administration, you can add, enable or disable rules related to Device Manager.

DHCP Rules Impacting Device Manager

The DHCP can interact directly with Device Manager thanks to two rules. These rules automatically create devices containing an interface every time you add a static or lease.

Note that the rules are first organized by modules and then event, so even if they both ultimately impact Device Manager, you will find them under the module *DHCP* and the event *Add: <DHCP-object>*.

Rule 221

Enabling this rule creates an interface every time you add a static, with or without an IPv4 address. The interface has the same name as the static and belongs to a device also named after the static. To add this rule, refer to the procedure [To add a Device Manager rule](#) and select the Module *DHCP* and the Event *Add: DHCP statics*.

Rule 225

Enabling this rule creates an interface every time an IPv4 lease is generated. The interface is named *generic_#* and belongs to a device named using the lease hostname. To add this rule, follow the procedure [To add a Device Manager rule](#) and select the Module *DHCP* and the Event *Add: DHCP leases*.

Adding Device Manager Rules

From the page *Rules*, you can add the rules of your choice. By default, none of Device Manager related rules are added.

To add a Device Manager rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the menu, click on  **Add**. The wizard **Add a rule** opens.
4. In the drop-down list **Module**, select the module that triggers the needed behavior. The rules impacting Device Manager are listed in the section [DHCP Rules Impacting Device Manager](#).
5. In the drop-down list **Event**, select the action in the selected module that triggers the behavior.
6. In the list **Rule**, select the rule of your choice. Each rule is listed as follows: (*<rule-number>*) *<rule-name>*.
7. In the **Rule name**, name the rule. This name is displayed in the column *Instance* and help you filter the list without using the rule number.
8. In the field **Comment**, you can type in a comment.
9. Click on **NEXT**. The last page of the wizard opens.

10. Click on to complete the operation. The report opens and closes. The rule is listed and marked  **OK** in the column **Status**.

Enabling or Disabling Device Manager Rules

At any time, you can enable or disable a rule, only the enabled rules are executed

To enable/disable Device Manager rules

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the search engine of the column **Rule #**, type in the partial or complete rule(s) number and hit Enter to filter the list.
4. Tick the rule(s) of your choice.
5. In the menu, select  **Edit** > **Enable** or **Disable**. The wizard opens.
6. Click on to complete the operation. The report opens and closes. In the column **Status**, the rule is marked  **OK** or  **Disabled**.

Part XIV. VLAN Manager

VLAN Manager allows to create and handle Virtual Local Area Networks (VLANs) and Virtual Extensible LAN (VXLAN) to set up layer 2 data exchange between networks and devices.

You can connect VLANs and organize IPAM subnet-type networks belonging to different spaces or networks using MAC addresses, therefore any device and network can be linked no matter their IP address.

Keep in mind that the VLAN Manager VLANs are different from the VLAN interfaces you can set up on the page *Network configuration*. For more details, refer to the section [Setting up a VLAN Interface](#).

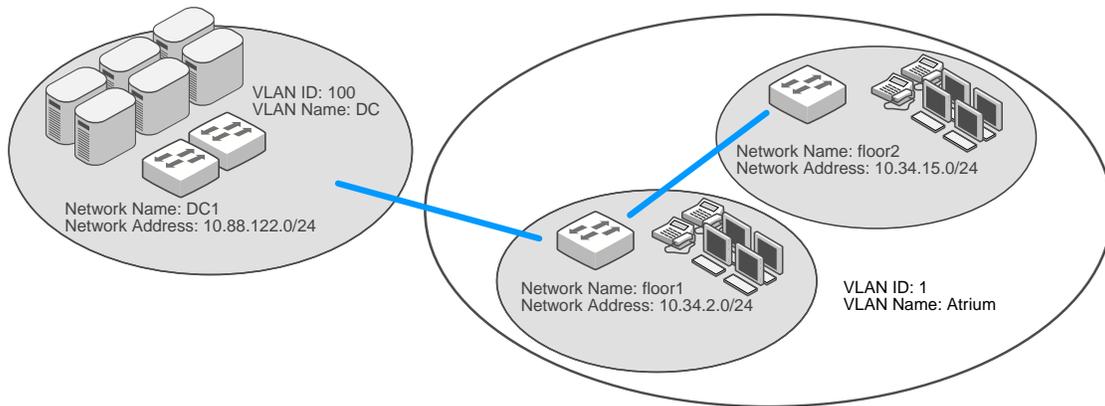


Figure 224. Example of a VLAN associating two subnet-type networks

From the module VLAN Manager you can include up to 3 levels of organization:

- **Domains:** the highest level of the hierarchy. They contain ranges and VLANs. For more details, refer to the chapter [Managing VLAN Domains](#).
- **Ranges:** an optional second level in the hierarchy. They contain VLANs. For more details, refer to the chapter [Managing VLAN Ranges](#).
- **VLANs:** the lowest level of the hierarchy. They are unique with a VLAN Identifier (ID) and belong to a domain or range. For more details, refer to the chapter [Managing VLANs](#).

Once you organized your VLANs, you can **set up the interaction with the IPAM**. For more details, refer to the chapter [Managing the IPAM/VLAN Interaction](#).

Note that from the module **Dashboards**, you can gather gadgets and charts on *VLAN Manager dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

Chapter 77. Managing VLAN Domains

VLAN domains are managed from the page **All domains**. They can be composed of VLAN ranges and VLANs or exclusively of VLANs depending on your organizational needs.

You need at least one domain to organize your VLANs.

A domain is defined by its name and start and end ID. These IDs corresponds to the first and last VLAN ID that it manages, it sets the number of VLANs it can contain:

- A VLAN domain can contain between 1 and 4094 VLANs.
- A VXLAN domain can contain between 1 and 16777215 VLANs.

Every time you add a domain, you can set the same set of IDs. They are duplicated on the page *All VLANs*, and even if you have several VLANs with the ID 1, they are different. Indeed, they do not belong to the same domain or range and might be assigned different names.

Browsing VLAN Domains

Within VLAN Manager, the domains are the highest level of the hierarchy. They are required to manage your VLANs.

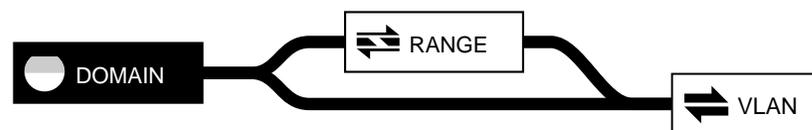


Figure 77.1. The domain in VLAN Manager hierarchy

Browsing the VLAN Domains Database

To display the list of VLAN domains

1. In the sidebar, go to **VLAN Manager > Domains**. The page **All domains** opens.
2. In the column **Domain End ID**, the values can help differentiate VLAN and VXLAN domains.

To display a VLAN domain properties page

1. In the sidebar, go to **VLAN Manager > Domains**. The page **All domains** opens.
2. At the end of the line of the domain of your choice, click on **ⓘ**. The properties page opens.

Customizing the Display on the Page All Domains

Users of the group *admin* can create customized column layouts. The button **☰ Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding VLAN Domains

Adding a domain sets the number of VLANs that you manage. You can set the start and end VLAN ID of your choice, for instance you can choose to manage the VLANs 25 to 500.

You can add as many domains as you want or import existing ones from a CSV file. For more details, refer to the section [Importing VLAN Domains](#).

Note that once you created a domain, you cannot edit its start and end ID, or decide to make it a VXLAN domain instead of VLAN or vice versa.

To add a VLAN domain

1. In the sidebar, go to **⇒ VLAN Manager > Domains**. The page **All domains** opens.
2. In the menu, click on **+ Add**. The wizard **Add a VLAN domain** opens.
3. If you or your administrator created classes at the domain level, in the list **Domain class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the field **Domain**, type in the VLAN domain name.
5. In the field **Description**, you can add a description. This field is optional. By default, *1* is displayed in the field.
6. To create a VXLAN domain, tick the box **Use VXLAN**.
7. In the field **Starting VLAN ID**, specify the ID of the first VLAN of the domain.
 - a. For a VLAN domain, type in a number between *1* and *4094*.
 - b. For a VXLAN domain, type in a number between *1* and *16777215*.
8. In the field **Ending VLAN ID**, specify the ID of the last VLAN of the domain.
 - a. For a VLAN domain, type in a number between *1* and *4094*. By default, *4094* is displayed in the field.
 - b. For a VXLAN domain, type in a number between *1* and *16777215*. By default, *16777215* is displayed in the field.
9. Click on **OK** to complete the operation. The report opens and closes. The domain is listed.

Editing VLAN Domains

Editing a domain means renaming it or changing its description, setting it, editing it or removing it.

You can edit a VLAN domain from the page *All domains*, via the contextual menu, or from its properties page.

If a domain no longer matches your needs and you want to edit its start ID, end ID, make it a VXLAN domain or make it a VLAN domain, you must:

1. Create a new domain configured with the settings that suit your needs.

You can either recreate the VLANs it contains and name them the same or export the VLANs of the domain you want to edit and reimport them into the domain you created. For more details, refer to the part [Imports and Exports](#).

2. Delete the obsolete domain.

To edit a VLAN domain

1. In the sidebar, go to **VLAN Manager > Domains**. The page **All domains** opens.
2. Right-click on the **Name** of the domain of your choice. The contextual menu opens.
3. Click on **Edit**. The wizard **Add a VLAN domain** opens.
4. If you or your administrator created classes at the domain level, in the list **Domain class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the fields **Domain** and/or **Description** according to your needs.
6. Click on **OK** to complete the operation. The report opens and closes. The domain is listed with the changes you just made.

Deleting VLAN Domains

Deleting a domain is only possible if it does not contain any range or *Used* VLANs. If any of its VLANs were assigned a name, you cannot delete the domain.

To delete a VLAN domain

1. In the sidebar, go to **VLAN Manager > Domains**. The page **All domains** opens.
2. Tick the domain(s) you want to delete.
3. In the menu, click on **Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The domain is no longer listed, the VLANs it contained are deleted from the page *All VLANs* as well.

Defining a VLAN Domain as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a VLAN domain as one of the resources of a specific group allows the users of that group to manage the VLAN domain in question as long as they have the corresponding rights granted.

Granting access to a domain as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 78. Managing VLAN Ranges

VLAN ranges provide an extra level of management for your VLANs. They are optional.

A VLAN range can contain as many VLANs as the domain it belongs to, or just a portion of the VLANs.

Like the domain, a range is defined by its name, its start ID and its end ID. Considering that it belongs to a domain, it cannot manage VLANs that are not managed by the domain, in other words you cannot create a range with the start and end IDs 5-10 in a domain managing the IDs 6-10.

Within a domain, you can create as many ranges as you want to manage the VLANs of the domain. Your ranges can manage the same VLAN IDs if you allow overlapping, the VLANs are different as they belong to different ranges.

Browsing VLAN Ranges

Within VLAN Manager, the ranges are the second level of the hierarchy. You can create ranges



Figure 78.1. The range in VLAN Manager hierarchy

Browsing the VLAN Ranges Database

To display the list of VLAN ranges

1. In the sidebar, go to **VLAN Manager > Ranges**. The page **All ranges** opens.
2. To display the list of ranges of a specific domain, at the end of the line of the range of your choice, click on **[icon]**. The properties page opens.
3. Click on the **Domain** name. The properties page of the domain is displayed.
4. In the breadcrumb, click on **All ranges**. The page **All ranges** of the selected domain opens.

To display a VLAN range properties page

1. In the sidebar, go to **VLAN Manager > Ranges**. The page **All ranges** opens.
2. At the end of the line of the range of your choice, click on **[icon]**. The properties page opens.

Customizing the Display on the Page All Ranges

Users of the group *admin* can create customized column layouts. The button **[icon] Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding VLAN Ranges

You can create as many ranges as you need in a domain.

If you want to create ranges with unique sets of VLAN ID, you can tick the box **No ID overlapping**. Keep in mind that the overlap restriction applies whether it was set on existing ranges or on ranges you are trying to create. Therefore, if a range managing the VLAN IDs *1-512* already exists and you try to create the range *512-550*, an error message is returned whether the box was ticked on the existing range or on the range you are creating.

With the overlapping allowed, if you set several ranges with common VLANs, the common VLAN ID is replicated on the page *All VLANs*. You can differentiate them through their range, and potentially their name.

Note that you can also import ranges from a CSV file. For more details, refer to the section [Importing VLAN Ranges](#).

To add a VLAN range

1. In the sidebar, go to **VLAN Manager > Ranges**. The page **All ranges** opens.
2. In the menu, click on **+ Add**. The wizard **Add a VLAN range** opens.
3. If you or your administrator created classes at the range level, in the list **Range class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

4. In the list **Domain**, select the domain of your choice.
5. Click on **NEXT**. The last page of the wizard opens.
6. In the field **Range**, name your VLAN range.
7. In the field **Description**, you can add a description. This field is optional.
8. In the field **Starting VLAN ID**, type in the ID of the first VLAN of the range. By default, the *Starting VLAN ID* of the domain is displayed in the field.
9. In the field **Ending VLAN ID**, type in the ID of the last VLAN of the range. By default, the *Ending VLAN ID* of the domain is displayed in the field.
10. The box **No ID overlapping** is ticked by default, you can untick it if you want to allow VLAN ID overlapping in several ranges.
11. Click on **OK** to complete the operation. The report opens and closes. The range is listed.

Editing VLAN Ranges

You can edit a range name and description or resize it, that is say manage more or less VLANs.

Editing a Range Properties

Editing a range means renaming it or changing its description, setting it, editing it or removing it.

You can edit a VLAN range from the page *All ranges*, via the contextual menu, or from its properties page.

To edit a VLAN range from the properties page

1. In the sidebar, go to **VLAN Manager > Ranges**. The page **All ranges** opens.
2. At the end of the line of the range of your choice, click on **ⓘ**. The properties page opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard **Add a VLAN range** opens.
4. If you or your administrator created classes at the range level, in the list **Range class**, select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Edit the fields **Domain**, **Description** according to your needs.
6. Tick or untick the box **No ID overlapping** according to your needs.
7. Click on **OK** to complete the operation. The report opens and closes. The range properties are updated.

Resizing a Range

At range level, you can change the number of VLANs managed. You can decide to manage more or less VLANs, i.e. VLAN IDs. This option basically shifts the VLAN identifier number to add IDs to the VLAN range or remove some IDs.

This option respects a set of rules:

1. You cannot reduce the size of a range if it contains *Used* VLANs (i.e. VLAN that were assigned a name and might therefore be linked to a subnet-type network in the IPAM)
2. You can extend the size of a range as much as you want provided that:
 - The new range size is not greater than the domain it belongs to.

In a domain managing the IDs 1-15, you cannot resize a range and make it manage the IDs 10-20, instead of 10-15. You would be asking to manage IDs that do not exist in the domain.

- The new range size does not exclude any *Used* VLAN of the range or include any *Used* VLAN belonging to another range.

In case of overlap, you can either delete the used VLAN and recreate it in the new range or export it and reimport it in the new range.

To resize a VLAN range

1. In the sidebar, go to **VLAN Manager > Ranges**. The page **All ranges** opens.
2. Tick the range(s) that you want to resize.
3. In the menu, select **⌵ Edit > Resize ranges**. The wizard **Resize ranges** opens.
4. In the **Start ID shift**, type in the value of your choice. The value can be positive or negative, preceded by -. A negative shift extends the number of IDs managed. If you do not want to edit the Start ID, type in 0.
5. In the **End ID shift**, type in the value of your choice. The value can be positive or negative, preceded by -. A negative shift reduces the number of IDs managed. If you do not want to edit the End ID, type in 0.

6. Click on **OK** to complete the operation. The report opens and closes. The new range(s) size is visible.

Deleting VLAN Ranges

Deleting a range is only possible if it does not contain any *Used* VLANs. If any of its VLANs were assigned a name, you cannot delete the range.

To delete a VLAN range

1. In the sidebar, go to **≡ VLAN Manager > Ranges**. The page **All ranges** opens.
2. Tick the range(s) you want to delete.
3. In the menu, click on **🗑 Delete**. The wizard **Delete** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The range is no longer listed, the VLANs it contained are deleted from the page *All VLANs* as well.

Defining a VLAN Range as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a VLAN range as one of the resources of a specific group allows the users of that group to manage the VLAN range in question as long as they have the corresponding rights granted.

Granting access to a range as a resource also grants access to every item it contains. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 79. Managing VLANs

Once you have created at least one domain, the VLANs it contains are listed on the page *All VLANs*. They can belong to ranges.

All the VLANs are differentiated through their ID. You can assign them a name to set up an interaction with the module IPAM at network level between several subnet-type networks or devices within a network. For this reason, once a VLAN has a name, the range and/or domain it belongs to cannot be deleted.

You can add, edit or delete VLANs.

Browsing VLANs

Within VLAN Manager, the VLANs are the lowest level of the hierarchy.



Figure 79.1. The VLAN in VLAN Manager hierarchy

Browsing the VLANs Database

To display the list of VLANs

1. In the sidebar, go to **VLAN Manager > VLANs**. The page **All VLANs** opens.

Only the *Used* VLANs are listed. For more details about statuses, refer to the section [Understanding the VLAN Statuses](#).

2. If the VLAN ID overlapping is enabled, the columns **Domain** and **Range** help differentiate VLAN IDs.

To display the list of VLANs of a specific VLAN domain

1. In the sidebar, go to **VLAN Manager > Domains**. The page **All Domains** opens.
2. In the column **Name**, click on the domain of your choice. The page **All VLANs** opens.

Only the VLANs of the selected domain are displayed, the *Used* ones and the first and last *Free* ones.

3. To display more *Free* VLANs, in the column **VLAN ID** click on **⊞** left of the ID of your choice. The first or last 13 VLANs appear.

To display a VLAN properties page

1. In the sidebar, go to **VLAN Manager > VLANs**. The page **All VLANs** opens.
2. At the end of the line of the *Used* VLAN of your choice, click on **ⓘ**. The properties page opens.

Customizing the Display on the Page All VLANs

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Understanding the VLAN Statuses

The column **Status** provides information regarding the VLANs you manage.

Table 79.1. VLAN statuses

Status	Description
 <i>Free</i>	The VLAN can be assigned a name.
 <i>Used</i>	The VLAN was assigned a name, it can interact with the IPAM. The domain and range it belongs to cannot be deleted.

Adding VLANs

Adding a VLAN means using it as all the VLANs were added at the same time as the domain they belong to. When you add a VLAN, you can assign it a name.

Note that you can also import VLANs from a CSV file. For more details, refer to the section [Importing VLANs](#).

To add a VLAN from the menu

1. In the sidebar, go to  **VLAN Manager** > **VLANs**. The page **All VLANs** opens.
2. In the menu, click on  **Add**. The wizard **Add a VLAN** opens.
3. In the list **Domain**, select the domain of your choice.
4. Click on **NEXT**. The next page of the wizard opens.
5. In the list **Range**, select the range of your choice or *none*.
6. Click on **NEXT**. The last page of the wizard opens.
7. In the field **VLAN name**, you can name the VLAN. This field is optional.
8. In the column **VLAN ID**, type in the VLAN ID of your choice.
9. Click on **OK** to complete the operation. The report opens and closes. The page refreshes and the VLAN is now marked as  *Used*. If you gave it a name, it is displayed in the column **Name**.

To add a VLAN from the listing page

1. In the sidebar, go to  **VLAN Manager** > **Domains** or **Ranges**. The page **All Domains** or **All Ranges** opens.
2. In the column **Name**, click on the domain or range of your choice. The page **All VLANs** opens.
3. Filter the list if need be.
4. In the column **VLAN ID**, click on the VLAN of your choice. A pop-up window **This VLAN ID is free, do you want to assign it?** opens.
5. Click on **OK**. The wizard **Add a VLAN** opens

6. In the field **VLAN name**, you can name the VLAN. This field is optional.
7. In the field **VLAN ID**, the VLAN ID you click on is displayed in gray.
8. Click on to complete the operation. The report opens and closes. The page refreshes and the VLAN is now marked as  *Used*. If you gave it a name, it is displayed in the column **Name**.

Editing VLANs

You can edit a VLAN name.

Keep in mind that renaming a VLAN breaks the IPAM / VLAN interaction.

To edit a VLAN

1. In the sidebar, go to  **VLAN Manager** > **VLANs**. The page **All VLANs** opens.
2. Filter the list if need be.
3. Right-click over the **VLAN ID** of the *Used* VLAN of your choice. The contextual menu opens.
4. Click on  **Edit**. The wizard **Add a VLAN** opens.
5. In the field **VLAN name**, rename the VLAN.
6. Click on to complete the operation. The report opens and closes. The new VLAN name is displayed.

Deleting VLANs

You can only delete *Used* VLANs. Before proceeding keep in mind that:

- Once deleted, the VLANs are still listed, if you are displaying the VLANs of a specific domain or range, but are *Free* and no longer have a name.
- If the VLANs were associated with a subnet-type network, deleting VLANs breaks the association and removes the VLAN information from the network properties. For more details, refer to the chapter [Managing the IPAM/VLAN Interaction](#).

The *Free* VLANs are unused so you cannot delete them. They are listed on the page *All VLANs* of the domain or range they belong to. To delete unused VLANs from the list, you must delete the range and/or domain they belong to.

To delete a VLAN

1. In the sidebar, go to  **VLAN Manager** > **VLANs**. The page **All VLANs** opens.
2. Tick the *Used* VLAN(s) of your choice.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on to complete the operation. The report opens and closes. The VLAN is no longer listed.
5. To see this VLAN, display the list of the VLANs for the domain or range it belongs to. This VLAN no longer has a name in the column **Name** and its status is  *Free*.

Chapter 80. Managing the IPAM/VLAN Interaction

The purpose of VLAN Manager is to create and control the interaction between virtual local area networks and your IPAM subnet-type networks.

Within the IPAM module, this interaction is managed at network level via advanced properties and can be set both from IPv4 and IPv6 networks.

Like any advanced property, you must display the fields in the relevant wizard and then configure them.

Displaying the IPAM/VLAN Interaction Advanced Properties

You can configure advanced properties to set up the IPAM/VLAN interaction. From the wizard *Advanced properties customization* you must configure which properties to display.

The first property allows associate a subnet-type network with an existing VLAN, the second property allows to add the VLAN on the page All VLANs when you associate it with a subnet-type network.

Table 80.1. IPAM/VLAN interaction advanced properties

In the wizard Advanced properties customization	In the addition/edition wizard
Display the VLAN association fields	VLAN domain
	VLAN range
	VLAN ID
Display the field "Create a VLAN"	Create a VLAN

To configure the IPAM/VLAN interaction advanced properties

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. In the menu, select **Extra options > Wizard customization**. The wizard **Advanced properties customization** opens.
4. In the section **IPAM/VLAN interaction**, tick the box **Display the VLAN association fields**. The box *Display the field "Create a VLAN"* appears.
5. Tick the box **Display the field "Create a VLAN"**.
6. Click on **OK** to complete the operation. The report opens and closes. The page **All networks** is visible again.
7. To set up the interaction in IPv4 and IPv6, follow the procedure again on the other version page.

Once you set the advanced properties display, you can configure the interaction.

Configuring the IPAM/VLAN Interaction

Once the properties are displayed, the interaction fields are available in the subnet-type networks addition/edition wizards in IPv4 or IPv6.

To make two networks communicate, you must configure them with a common VLAN. That configuration enables them to send/receive packets, etc. no matter what block-type network either network belongs to. **The network/VLAN interaction can be set between a terminal or a non-terminal subnet-type network and a VLAN.**

To link a subnet-type network to an existing VLAN

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit a subnet-type network. The corresponding wizard opens. For more details regarding the first steps of addition or edition, refer to the chapter [Managing Networks](#).

If you edit an existing non-terminal subnet-type network that contains terminal networks, linking it to a VLAN does not link the networks it contains to the VLAN.

4. On the last page of the wizard, in the drop-down list **VLAN domain**, select the VLAN domain containing the VLAN you want to associate with your subnet-type network.
5. In the drop-down list **VLAN range**, select the value of your choice:

Table 80.2. Available options in the drop-down list VLAN range

Option	Description
<range-name>	Select the name of an existing range to narrow down the auto-complete search of the field <i>VLAN ID</i> .
None	Select this option if there is no range in the domain or if the VLAN you are looking for is not managed by a range.
All	Select this option if you do not know in which range the VLAN you are looking for is. This option is not available if several ranges in your domain have the same VLAN ID assigned, even if they have a different name.

6. In the field **VLAN ID**, type in the first digit(s) of the VLAN ID you are looking for. The field auto-completes and displays the matching VLAN in the field or provides a list of the matching VLAN IDs (partially or entirely).

The VLAN is displayed as follows: <VLAN_ID> (<VLAN_name> - <range_name>). Where <range_name> can be replaced by #, if there is no range.

7. Click on **OK** to complete the operation. The report opens and closes. The VLAN configuration is visible on the network properties page in the panel **Advanced properties**.

To create and link a VLAN to a subnet-type network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. Add or edit a subnet-type network. For more details, refer to the chapter [Managing Networks](#). The corresponding wizard opens.

4. On the last page of the wizard, in the drop-down list **VLAN domain**, select the VLAN domain managing the VLAN you want to associate with your subnet-type network.
5. In the drop-down list **VLAN range**, select the value of your choice:

Table 80.3. Available options in the drop-down list VLAN range

Option	Description
<range-name>	Select the name of an existing range to narrow down the auto-complete search of the field <i>VLAN ID</i> .
None	Select this option if there is no range in the domain or if the VLAN you are looking for is not managed by a range.
All	Select this option if you do not know in which range the VLAN you are looking for is. This option is not available if several ranges in your domain have the same VLAN ID assigned, even if they have a different name.

6. Tick the box **Create a VLAN**. The field *VLAN name* appears.
7. In the field **VLAN ID**, type in the VLAN ID of the VLAN you want to create or type in the first digit(s) of the VLAN ID you are looking for. The field auto-completes and displays all the available IDs in the domain or range you selected.

The VLAN ID is displayed as follows: <VLAN_ID> (<range_name>). Where <range_name> can be replaced by #, if the VLAN is not managed through a range.

8. In the field **VLAN name**, you can name the VLAN.
9. Click on **OK** to complete the operation. The report opens and closes. The VLAN configuration is visible on the network properties page in the panel **Advanced properties**.

Once the association is set, you can display its details on the page *All networks* in the columns **VLAN Domain**, **VLAN Range**, **VLAN name** and **VLAN ID**. For more details, refer to the section [Customizing the List Layout](#).

Removing the IPAM/VLAN Interaction

At any time, you can remove the IPAM/VLAN interaction. There are two ways of proceeding:

1. Removing the association of one subnet-type network with a VLAN from the IPAM module.
2. Removing the association of a VLAN with all its subnet-type networks from the module VLAN Manager.

To remove the link between one subnet-type network and a VLAN

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4** or **V6** depending on your needs. The page refreshes and the button turns black.
3. Right-click on the **Name** of the subnet-type network of your choice. The contextual menu opens.
4. Click on **Edit**. The wizard opens.
5. If you or your administrator created classes, the page **Network class** opens.
6. Click on **NEXT**. The last page opens.

7. In the drop-down list **VLAN domain**, select *None*. The fields **VLAN range** and **VLAN ID** are no longer visible.
8. Click on to complete the operation. The report opens and closes. On the network properties page, in the panel **Advanced properties**, the VLAN configuration details are set to *None*.

To remove a VLAN association with all subnet-type networks at once

1. In the sidebar, go to  **VLAN Manager** > **VLANS**. The page **All VLANS** opens.
2. Tick the VLAN(s) of your choice.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on to complete the operation. The report opens and closes. The VLAN is listed but no longer has a name in the **Name** column and its status is now  *Free*. On the properties page of all the networks that were associated with the VLAN(s), in the panel **Advanced properties**, the VLAN configuration details are set to *None*.

Part XV. VRF

VRF, Virtual Routing and Forwarding, allows to simultaneously define and maintain multiple instances of a routing table on a single router.

This technology is commonly used for implementing L3 VPN(s) provided by MPLS service providers. In such networks MPLS encapsulation is used to isolate individual customer traffic, and an independent routing table (VRF) is maintained for each one of them.

Following RFC 4364, each VRF has a unique Route Distinguisher (RD) identifier. In that context, MP-BGP is commonly employed to facilitate complex redistribution schemes to import and export routes to and from VRFs (using route targets) to provide Internet connectivity or inter-VRF communication. Technically, you should keep in mind that:

- Each VRF behaves like an independent router with its own interfaces, IP subnet-type networks and routing protocol.
- Each VRF has separate routing and forwarding tables used only for the packets that traffic through the VRF based on its interface mapping.

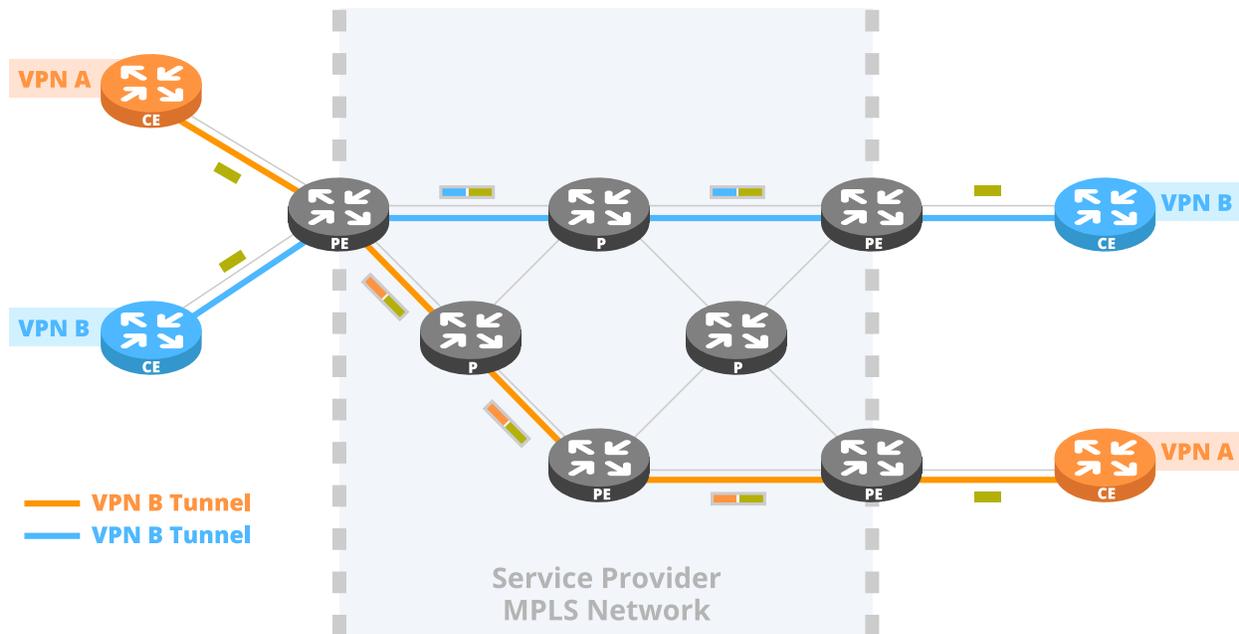


Figure 228. Example of use of VRF on an MPLS network

From the module **VRF**, you can display and have an overview of the VRF and Route Targets that associate them on your network on two dedicated pages. All available options are detailed in the chapters:

- [Managing Virtual Routing and Forwarding.](#)
- [Managing VRF Route Targets.](#)

Chapter 81. Managing Virtual Routing and Forwarding

From the page *All VRFs*, you can add, import, edit and delete Virtual Routing and Forwarding (VRF) for basic management purposes.

This page inventories all the VRFs you have on your network using their name and unique Route Distinguisher or RD.

When all your VRFs are listed, you can oversee the communication configurations between them via the Route Targets. For more details, refer to the chapter [Managing VRF Route Targets](#).

Browsing VRFs

Within the module *VRF*, the VRFs are the highest level, the entry point of your inventory.

Browsing the VRF Database

To display the list of VRFs

1. In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
2. You can filter the list using the page columns.

To display a VRF properties page

1. In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
2. At the end of the line of the VRF of your choice, click on . The properties page opens.

Customizing the Display on the Page All VRFs

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding VRFs

You can add as many VRFs as you need.

Note that you can also import VRFs, for more details refer to the section [Importing Data to VRF](#).

To add a VRF

1. In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
2. In the menu, click on  **Add**. The wizard **Add a VRF** opens.
3. If you or your administrator created classes at the VRF level, in the list **VRF class** select a class or *None*. Click on **NEXT**. The page **Add a VRF** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- In the field **Name**, name your VRF.
- In the field **RD**, type in the Route Distinguisher of your VRF. Following RFC 4364, the three accepted types rule the RD format:

Table 81.1. RD types and formats

Type	RD Format
0	<integer between 0 and 65535>:<integer between 0 and 4294967296>
1	<IPv4 address>:<integer between 0 and 65535>
2	<integer between 0 and 4294967296>:<integer between 0 and 65535>

- In the field **Comment**, you can type in a comment. This field is optional.
- Click on to complete the operation. The report opens and closes. The VRF is listed.

Editing VRFs

Once created, you can edit all the information regarding a VRF, from the listing page or from its properties page in the panel *Main properties*.

Keep in mind that **editing a VRF name or RD also edits its VRF Route Target(s)** in the GUI.

To edit a VRF

- In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
- Right-click over the **Name** of the VRF you want to edit. The contextual menu opens.
- Click on . The wizard opens.
- If you or your administrator created classes at the VRF level, in the list **VRF class** select a class or *None*. Click on . The page **Add a VRF** opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

- Edit the **Name**, **RD** and **Comment** fields according to your needs.
- Click on to complete the operation. The report opens and closes. The page refreshes, the VRF is listed with the new information.

Deleting VRFs

At any point you can delete one or several VRFs.

Keep in mind that **deleting a VRF also deletes its VRF Route Targets** in the GUI.

To delete a VRF

- In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
- Tick the VRF(s) you want to delete.

3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on  to complete the operation. The report opens and closes. The VRF is no longer listed, its VRF Route Targets are deleted as well.

Chapter 82. Managing VRF Route Targets

A Route Target sets up an exchange of routes between two VRFs, via their RD. They allow to import and/or export the routes of the VRFs of your choice, one is the source VRF, the other is the target VRF.

From the page *All VRF Route Targets*, you can have an overview of the communication you set or want to set between the VRFs. Every VRF can be associated with one or more Route Targets.

The page simply illustrates the VRF communication configuration you set on your network and assists you in inventorying it all.

Browsing VRF Route Targets

Within the module *VRF*, the VRF Route Targets are the lowest level, they link together the VRFs.

Browsing the VRF Route Targets Database

To display the list of VRF Route Targets

1. In the sidebar, go to  **VRF > VRF Route Targets**. The page **All VRF Route Targets** opens.
2. You can filter the list using the page columns.

VRF Route Targets do not have a properties page as all the information is displayed on the page.

To display the list of VRF Route Targets of a specific source VRF

1. In the sidebar, go to  **VRF > VRFs**. The page **All VRFs** opens.
2. Click on the name of the VRF of your choice. The page **All VRF Route Targets** opens. Only the Route Targets that defined the VRF RD as *Source VRF* are listed.

Customizing the Display on the Page All VRF Route Targets

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding VRF Route Targets

From the page *All VRF Route Targets*, you can add Route Targets to inventory the ones that are already configured on your network or that you want to set up.

To properly add Route Targets you must:

- Have at least two VRFs listed on the page *All VRFs*.
- Specify a Source VRF and a Target VRF.
- Specify the exchange medium between the VRFs, importation and/or exportation:
 - **import**: the Target VRF can import all the routes of the Source VRF.
 - **export**: the Source VRF sends out its routes to the Target VRF.

- Make sure that you create two Route Targets to establish and confirm your exchange configuration. For instance:
 1. Create a Route Target that establishes that the Source VRF *site A* can *import* the routes of the Target VRF *site B*.
 2. Create another Route Target that establishes that the Source VRF *site B* can *export* the routes the Target VRF *site A*.

To properly set up the exchange of routes, an *import* must be confirmed by an *export*, and vice versa.

Note that:

- You can also import Route Targets, for more details refer to the section [Importing VRF Route Targets](#).
- A Route Target may be edited if its Source or Target VRF is edited.
- You cannot edit Route Targets from the page *All VRF Route Targets*. You must delete the Route Target that no longer suits your needs and create it again.

To add the Route Targets that establish the communication between two VRFs

1. In the sidebar, go to  **VRF > VRF Route Targets**. The page **All VRF Route Targets** opens.
2. **Create the first Route Target**
 - a. In the menu, click on **+ Add**. The wizard **Add a VRF Route Target** opens.
 - b. In the field **Source VRF name**, specify the VRF of your choice.

Type in the first letters of the source VRF, the auto-completion provides the list of matching names, select the one you want.
 - c. In the field **Target VRF name**, specify the VRF of your choice using auto-completion.
 - d. You can tick the box **Import** to let the Target VRF retrieve the routes of the Source VRF.
 - e. You can tick the box **Export** to let the Source VRF send out its routes to the Target VRF.
 - f. Click on to complete the operation. The report opens and closes. The VRF is listed.
3. **Create the second Route Target**
 - a. In the menu, click on **+ Add**. The wizard **Add a VRF Route Target** opens.
 - b. In the field **Source VRF name**, specify the Target VRF of the first Route Target.
 - c. In the field **Target VRF name**, specify the Source VRF of the first Route Target.
 - d. If you ticked the box *Export* in the first Route Target, tick the box **Import** to confirm the communication configuration.
 - e. If you ticked the box *Import* in the first Route Target, tick the box **Export** to confirm the communication configuration.
 - f. Click on to complete the operation. The report opens and closes. The VRF is listed.

Deleting VRF Route Targets

At any point you can delete a VRF Route Target between two VRFs.

Note that Route Targets can be deleted from the page when one of the VRFs they link is deleted.

To delete a VRF Route Target

1. In the sidebar, go to  **VRF > VRF Route Targets**. The page **All VRF Route Targets** opens.
2. Tick the VRF Route Target(s) you want to delete.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on to complete the operation. The report opens and closes. The page refreshes, the VRF Route Target is no longer listed.

Part XVI. SPX

The Service Provider eXtension (SPX) assists Local Internet Registry (LIR) declarations as it allows to manage the complete life cycle of the IP address networks allocated to you by a Regional Internet Registry (RIR) member. From SOLIDserver GUI, SPX helps you manage networks that were allocated to you by the RIPE (Réseaux IP Européens) or the APNIC (Asia-Pacific Network Information Center).

The module SPX comes in addition to the IPAM and is available through a dedicated license option. To make sure you do have this license option, the administrator can go to the page *Admin Home* and, in the section *System*, click on *License*. When the page opens, in the panel *Current license*, all the license options are listed: *SPX* must be listed.

To properly use SPX you must:

1. Configure SOLIDserver with the organization details of your RIR, whether the RIPE or APNIC. For more details, refer to the chapter [Configuring SPX](#).
 2. Add or import the user(s), i.e. RIPE or APNIC person(s), responsible for the SPX network(s) management. For more details, refer to the chapter [Managing SPX Persons](#).
 3. Import the networks that your RIR allocated to you and add or import the assigned network that suit your needs. For more details, refer to the chapter [Managing SPX Networks](#).
 4. Add or import the AS Numbers registration details, i.e. RIPE or APNIC aut-nums, that suit your needs. For more details, refer to the chapter [Managing SPX AS Numbers](#).
-

Chapter 83. Configuring SPX

No matter what RIR you depend on, there is only one wizard to configure SOLIDserver. Once SPX is properly set and matches your allocated network(s), only assigned networks can be added and edited: their containers are managed by the RIR itself. Whenever you add or edit assigned networks from the GUI, an email is sent to your RIR.

Enabling the SPX Classes

To properly configure SOLIDserver to manage RIR networks, you need to enable the default SPX classes provided. They allow to add the extra fields and options that assist you in managing RIPE or APNIC networks.

To enable SPX classes

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the search engine of the column **Directory**, type in *SPX*. Only the default SPX classes are listed.
4. Next to the column **Name** tick the box: all the classes of the directory are selected.
5. In the menu, select  **Edit** > **Enable class**. The wizard **Enable class** opens.
6. Click on to complete the operation. The report opens and closes, the page refreshes. The classes are marked as  *Enabled* in the column **Status**.

Enabling the SPX Rules

Once the SPX classes are enabled, you must enable the default SPX rules. These rules are designed to automate the communications between SOLIDserver and the RIPE or APNIC organization.

To enable SPX rules

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the search engine of the column **Module**, type in *SPX*. The SPX rules are listed.
4. Next to the column **Name** tick the box: all the classes of the module are selected.
5. In the menu, select  **Edit** > **Enable**. The wizard **Enable** opens.
6. Click on to complete the operation. The report opens and closes, the page refreshes. The rules are marked as  *OK* in the column **Status**.

Configuring the SPX Connection

To configure SOLIDserver with your RIR organization details and set up management preferences using SOLIDserver classes, a configuration wizard is available in the module SPX.

Before going further, make sure that:

- You have all your RIPE or APNIC network details: maintainer, organization, registry identifier, administrator contact (admin-c) and user contact (person).
- Your SPX classes and rules are enabled. They apply to IPv4 and IPv6 allocated networks, IPv4 and IPv6 assigned networks, autnums and finally users. For more details, refer to the sections [Enabling the SPX Classes](#) and [Enabling the SPX Rules](#).

Keep in mind that this configuration wizard allows to configure your RIPE or APNIC "real" database as well as your "TEST" database, if you have one.

To configure SOLIDserver SPX with your RIPE or APNIC details

1. In the sidebar, go to  **SPX > AS Numbers or Policies**. The page opens.
2. In the menu, select  **Tools > SPX configuration**. The wizard **SPX configuration** opens.
3. Configure your SPX settings:

Field	Description
RIR	Select either <i>RIPE</i> or <i>APNIC</i> .
Comment	Type in a comment regarding the organization.
Maintainer	Type in your RIPE or APNIC maintainer full name. This information is contained in the field <i>mntner</i> and reused in the field <i>mnt-by</i> of your assigned networks.
Password	Type in your RIPE or APNIC password. It is used to authenticate the database updates.
Source	Select either the RIR you selected above to configure your official database or <i>TEST</i> to configure your test database.
NCC REGID	Type in your registry identifier. It was provided to you by your RIR, if not, you should contact them to obtain it.
From (email)	Type in the email address used as source address in the emails sent to your RIR.
Notify (email)	Type in the email address of the person notified of any change made in the RIPE or APNIC database.
Changed (email)	Type in the email address displayed in the field <i>changed</i> of the description of your assigned network in the RIPE or APNIC database. It can be a generic address or the address of a person.
AW validation (email)	Type in the email of the person notified if you exceed the number of IP addresses of your Assigned Window. This person must be granted sufficient rights to perform the appropriate operations if your new assigned networks exceed the allocated range of addresses.
AW size	Type in the number of IP addresses you are allocated in the RIPE or APNIC Assigned Window.
Update method	Select <i>POST</i> or <i>EMAIL</i> .
<i>POST</i>	This service based method is selected by default to notify your RIR of any changes.
<i>EMAIL</i>	This method allows to notify your RIR of any changes via email. Once selected, the following fields appear.
	Update pop3 mailbox Type in the pop3 address of your mail server.
	Update mailbox login Type in the login of the specified mail server.
	Update mailbox password Type in the password of the specified mail server.
Expert mode	Tick this box if you want to set up a proxy server to communicate changes to your RIR. Once ticked, the following fields appear.
	Whois RIR host Type in the full name of the proxy server.

Field	Description	
	Whois port	Type in the number of the Whois RIR host port used to transmit information. Port 80 is generally used.
	RIR Update host	Type in the name of the server of your RIR receiving the updates
	RIR update URL	Type in the URL of the RIPE or APNIC server receiving your updates.
	Email used for the update	Type in the email address used as source address in the emails notifying your RIR of any update.

Once all the fields are filled, click on **ADD**.

The details are moved to the **Maintainer list** and displayed as follows: *Source: <selected-source> - Maintainer : <maintainer-name>*.

4. Repeat the configuration step for as many maintainers as needed.
5. Click on **NEXT**. The next page opens.
6. On the page **SPX allocated networks classes configuration**, configure the classes of your RIR allocated networks:
 - a. In the drop-down list **Allocated network class**, select one of your classes or the default class *SPX/RIPE_block*¹.
 - b. In the drop-down list **Allocated network PI class**, select one of your classes or the default class *SPX/RIPE_PI_block*.
 - c. Click on **NEXT**. The next page opens.
7. On the page **SPX assigned networks classes configuration**, configure the classes of your RIR assigned networks:
 - a. In the drop-down list **Assigned network class**, select one of your classes or on of the default classes, *SPX/RIPE_PI_subnet* or *SPX/RIPE_subnet*.
The selected class is moved to the field **New Assigned network class**.
 - b. Click on **ADD** to confirm the selection. The class is moved to the **List of SPX assigned networks**.
 - c. Repeat these steps for as many classes as needed.
 - d. Click on **NEXT**. The next page opens.
8. On the page **SPX allocated networks (v6) classes configuration**, configure the classes of your IPv6 RIR allocated networks:
 - a. In the drop-down list **Allocated network class (v6)**, select one of your classes or the default class *SPX/RIPE_Block*.
 - b. Click on **NEXT**. The next page opens.
9. On the page **SPX assigned networks (v6) classes configuration**, you can configure the classes of your IPv6 RIR assigned networks:

¹All the classes name can be preceded by a / if they belong to a specific directory, following the format: *<directory-name>/<class-name>*. In this case, the default class *RIPE_Block* belongs to the directory *SPX*.

- a. In the drop-down list **Assigned network class (v6)**, select one of your classes or the default class *SPX/RIPE_subnet*.
The selected class is moved to the field **New assigned network class (v6)**.
 - b. Click on **ADD** to confirm its selection. The class is moved to the **List of SPX assigned networks (v6)**.
 - c. Repeat the steps above for as many classes as needed.
 - d. Click on **NEXT**. The next page opens.
10. On the page **SPX aut-nums classes configuration**, you can configure the classes for your autnums:
- a. In the drop-down list **Aut-num class**, select one of your classes or the default class *SPX/RIPE*. The selected class is moved to the field **New aut-num class**.
 - b. Click on **ADD** to confirm its selection. The class is moved to the **List of SPX aut-num classes**.
 - c. Repeat these steps for as many classes as needed.
 - d. Click on **NEXT**. The next page opens.
11. On the page **SPX users classes configuration**, you can configure the classes of your RIPE or APNIC users:
- a. In the drop-down list **User class**, select one of your classes or the default class *SPX/RIPE_person*.
The selected class is moved to the field **New user class**.
 - b. Click on **ADD** to confirm its selection. The class is moved to the **List of SPX users**.
 - c. Repeat these steps for as many classes as needed.
12. Click on **OK** to complete the operation. The report opens and closes, the page refreshes.

At any time, you can edit these settings or add new maintainers. For more details, refer to the section [Editing the Connection to the RIPE or APNIC](#).

Editing the Connection to the RIPE or APNIC

Once your configuration with SOLIDserver is done, you can always edit its details or the class associated with your maintainer from the wizard *SPX configuration*.

Keep in mind that **you should not edit the maintainer name, registry identifier or AW size if you already imported your allocated networks**.

To edit an SPX maintainer configuration details

1. In the sidebar, go to  **SPX > AS Numbers or Policies**. The page opens.
2. In the menu, select  **Tools > SPX configuration**. The wizard **SPX configuration** opens.
3. At the bottom of the page in the **Maintainer list**, click on the maintainer you want to edit. The configuration current values are displayed in the each field.

4. Change the value of the field(s) of your choice. For more details regarding the fields, refer to the procedure in the section [Configuring the SPX Connection](#).
5. Click on **UPDATE**. The **Maintainer list** is edited according to your changes. Only the *Source* and *Maintainer* name are displayed on this list.
6. Click on **NEXT**. The page **SPX allocated networks classes configuration** opens.
7. Click on **NEXT**. The page **SPX assigned networks classes configuration** opens.
8. Click on **NEXT**. The page **SPX allocated networks (v6) classes configuration** opens.
9. Click on **NEXT**. The page **SPX assigned networks (v6) classes configuration** opens.
10. Click on **NEXT**. The page **SPX aut-num classes configuration** opens.
11. Click on **NEXT**. The page **SPX users classes configuration** opens.
12. Click on **OK** to complete the operation. The report opens and closes, the page refreshes.

To delete an SPX maintainer

1. In the sidebar, go to **SPX > AS Numbers or Policies**. The page opens.
2. In the menu, select **Tools > SPX configuration**. The wizard **SPX configuration** opens.
3. At the bottom of the page in the **Maintainer list**, click on the maintainer you want to delete. The configuration current values are displayed in the each field.
4. Click on **DELETE**. The maintainer is no longer in the **Maintainer list**.
5. Click on **NEXT**. The page **SPX allocated networks classes configuration** opens.
6. Click on **NEXT**. The page **SPX assigned networks classes configuration** opens.
7. Click on **NEXT**. The page **SPX allocated networks (v6) classes configuration** opens.
8. Click on **NEXT**. The page **SPX assigned networks (v6) classes configuration** opens.
9. Click on **NEXT**. The page **SPX aut-num classes configuration** opens.
10. Click on **NEXT**. The page **SPX users classes configuration** opens.
11. Click on **OK** to complete the operation. The report opens and closes, the page refreshes.

To edit the classes associated with an SPX maintainer

1. In the sidebar, go to **SPX > AS Numbers or Policies**. The page opens.
2. In the menu, select **Tools > SPX configuration**. The wizard **SPX configuration** opens.
3. At the bottom of the page in the **Maintainer list**, click on the maintainer which classes you want to edit. The configuration current values are displayed in the each field.
4. Click on **NEXT**. The page **SPX allocated networks classes configuration** opens.
5. You can select a different class in the drop-down lists **Allocated network class** and **Allocated network PI class**.
6. Click on **NEXT**. The page **SPX assigned networks classes configuration** opens.
7. You can edit the list of classes in the **List of SPX assigned networks classes**: select a class in the drop-down list **Assigned network class** and **ADD** it, or select a class in the list and **DELETE** it.
8. Click on **NEXT**. The page **SPX allocated networks (v6) classes configuration** opens.
9. You can select a different class in the drop-down list **Allocated network class (v6)**.

10. Click on **NEXT**. The page **SPX assigned networks (v6) classes configuration** opens.
11. You can edit the list of classes in the **List of SPX assigned networks classes (v6)**: select a class in the drop-down list **Assigned network class (v6)** and **ADD** it, or select a class in the list and **DELETE** it.
12. Click on **NEXT**. The page **SPX aut-num classes configuration** opens.
13. You can edit the list of classes in the **List of SPX aut-num classes**: select a class in the drop-down list **Aut-num class** and **ADD** it, or select a class in the list and **DELETE** it.
14. Click on **NEXT**. The page **SPX users classes configuration** opens.
15. You can edit the list of classes in the **List of SPX users classes**: select a class in the drop-down list **User class** and **ADD** it, or select a class in the list and **DELETE** it.
16. Click on **OK** to complete the operation. The report opens and closes, the page refreshes.

Chapter 84. Managing SPX Persons

To manage SPX networks, you first need to add or import RIPE and APNIC persons. From the GUI, SPX persons are considered as SOLIDserver users that you can manage as regular users for which you need to grant rights over the resources that suit your needs. For more details, refer to the part [Rights Management](#).

Browsing SPX Persons

SPX persons are visible in the module Administration on the page Users.

To display the list of SPX persons

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. On the right-end side of the menu, click on  **Listing templates**. The window opens.
4. To display the SPX persons related columns on the page, make sure the parameters described in the table below are ticked:

Table 84.1. SPX persons class parameters

Parameter	Column name
User class name	Class
Waiting	Waiting
Class param: Address	Address
Class param: Email	Email
Class param: Fax	Fax
Class param: Maintainer	Maintainer
Class param: Notify	Notify
Class param: Person	Person
Class param: Phone	Phone
Class param: Remark	Remark

Note that most of these parameters are only available if you correctly enabled the classes located in the directory SPX. For more details, refer to the chapter [Enabling the SPX Classes](#).

5. Click on . The page refreshes.
6. You can filter the list to display only SPX users by typing *RIPE* in the search engine of the column **Class**.

Adding SPX Persons

From the page *Users*, you can add SPX persons.

Once you added a person via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the person addition was confirmed, you can display the column **Waiting**. For more details, refer to the section [Browsing the SPX Persons](#).

To add a SPX person

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. In the menu, click on **+ Add**. The wizard **Add a user** opens.
4. In the list **User class**, select the class *RIPE_person* .
5. Click on **[NEXT]**. The page **Add a user** opens.
6. Configuring the SPX person details:
 - a. In the field **Usr login**, an identifier is automatically incremented. You can edit it if need be.
 - b. In the field **Address**, type in the person mailing address to fill in the RIPE or APNIC field *address*.
 - c. In the field **Phone**, type in the person phone number following the format: *+<country code> <area code> <phone number>*.
 - d. In the field **Fax**, you can type in a fax number following the same format as the field *Phone*.
 - e. In the field **Email**, type in the user email address.
 - f. In the field **Remark**, you can type in a comment regarding the person.
 - g. In the field **Notify**, you can type in the email address of the person notified of any changes made on the details of the person you are creating.
 - h. In the drop-down list **Maintainer**, select your maintainer.
7. If need be, configure extra details for the person management from the GUI:
 - a. In the field **First name**, type in the person first name.
 - b. In the field **Last name**, type in the person last name.
 - c. In the field **Pseudonym**, the user last and first name are automatically displayed. You can replace them by a shortname or shorter name if you want.
8. Click on **[OK]** to complete the operation. The report opens and closes. The user is listed, its state is  *Creating*. Until its status is not  *OK*, the RIPE or APNIC has not confirmed the addition. Have a look in the column **Waiting state** for more details regarding the addition confirmation.

If the person status stays in *wait_mail_add*, refer to the section [Registering SPX Person Changes](#).

Editing SPX Persons

SPX persons can be edited from the Users page. Any changes are sent to the RIPE or APNIC and the email address of the Notify field if it was set during the person creation.

Once you edited a person via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the assigned network edition was confirmed, you can display the column **Waiting**. For more details, refer to the section [Browsing the SPX Persons](#).

To edit a SPX person

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. In the column **Login**, click on the user name of your choice. The properties page opens.
4. In the panel **Main properties**, click on . The wizard **Edit a user** opens.
5. In the list **User class**, edit the class if need be.
6. Click on . The page **Edit a user** opens.
7. Edit the user information according to your needs. You can edit any detail except their *Pseudonym*, all parameters are detailed in the procedure [Adding SPX Persons](#).
8. Click on to complete the operation. The report opens and closes. The changes are listed in the panel.
9. Go back to the Users list to see the person state and make sure it was confirmed by the RIPE or APNIC. Until its status is not  **OK**, the RIPE or APNIC has not confirmed the edition. Have a look in the column **Waiting** for more details regarding the edition confirmation.

If the person status stays in *wait_mail_add*, refer to the section [Registering SPX Person Changes](#).

Registering SPX Person Changes

If the persons you added, edited or deleted have the status *wait_mail_add*, *wait_mail_del* or *must_send_mail*, you need to use the option **Register again**. This option resends your assigned network information to the RIPE or APNIC via POST or EMAIL, depending on your configuration.

This option resends your assigned network or person information to the RIPE or APNIC via POST or EMAIL, depending on your configuration.

To register again newly added or edited SPX assigned persons

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. Tick the person(s) that have the status *wait_mail_add* or *wait_mail_del*.
4. In the menu, select  **Edit > Register again**. The wizard **Person Register again** opens.
5. Click on to complete the operation. The report opens and closes. The assigned network Status evolves until it is  **OK**. Have a look in the column **Waiting state** for more details regarding the RIPE or APNIC confirmation.

Deleting SPX Persons

SPX persons can be deleted from the Users page. This deletion request is sent to the RIPE or APNIC and the email address of the Notify field if it was set during the person creation.

Once you deleted a person via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the assigned network edition was confirmed, you can display the column **Waiting**. For more details, refer to the section [Browsing the SPX Persons](#).

Before deleting a person, make sure that the assigned networks they were managing are already managed by someone else: edit the assigned networks concerned *Contacts* details. For more details, refer to the section [Editing SPX Networks](#).

To edit a SPX person

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. Tick the user(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The user state is  *Deleting* until the RIPE or APNIC confirms its deletion. Have a look in the column **Waiting state** for more details regarding the deletion confirmation.

If the person status stays in *wait_mail_del*, refer to the section [Registering SPX Person Changes](#).

Chapter 85. Managing SPX Networks

From SOLIDserver GUI, you can import allocated networks and edit their content by adding, editing and deleting assigned networks.

SPX assigned networks addition is assisted by dedicated classes that the administrator should enable in Class Studio, all these classes belong to the *Directory* SPX. For more details, refer to the chapter [Configuring SPX](#).

More generally, SPX networks can be managed the same way as regular networks, For more details, refer to the chapter [Managing Networks](#), [Managing Pools](#) and [Managing IP Addresses](#).

Note that SOLIDserver supports Provider Aggregatable and Provider Independent addresses¹. You can import or add them using the dedicated classes available for allocated networks and assigned networks.

Browsing SPX Networks

SPX networks are visible in the module IPAM, on the page All networks, where allocated networks are displayed as 🏠 block-type networks and assigned networks as 🏡 subnet-type networks.

Any addition sends a request to the RIPE or APNIC that is confirmed or denied. The status of that request can be displayed in a dedicated column on the page *All networks* in the module IPAM.

To display the list of SPX networks

1. In the sidebar, go to 🏡 **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on v4 or v6 depending on your needs. The page refreshes and the button turns black.
3. On the right-end side of the menu, click on **Listing templates**. The window opens.
4. To display the SPX networks related columns on the page, make sure the parameters described in the table below are ticked:

Table 85.1. SPX networks class parameters

Parameter	Column name
Network class name	Class
Network waiting state	Waiting state
* Description (__eip_description)	Description
* Description (descr)	Description
Class param: Admin c	Administrative contacts
Class param: Assigned network	Assigned network
Class param: Changed	Changed
Class param: Country	Country
Class param: Maintainer	Maintainer
Class param: Name	Name
Class param: Notify	Notify

¹For more details, refer to <http://www.ripe.net/lir-services/member-support/info/faqs/isp-related-questions/pa-pi>.

Parameter	Column name
Class param: Remarks	Remarks
Class param: Rev srv	Rev srv
Class param: Status	Status
Class param: Tech c	Technical contacts

Note that most of these parameters are only available if you correctly enabled the classes located in the directory SPX. For more details, refer to the chapter [Enabling the SPX Classes](#).

5. Click on **SAVE**. The page refreshes.
6. You can filter the list to display only SPX networks by typing *RIPE* in the search engine of the column **Class**.

Adding SPX Networks

From SOLIDserver, you cannot add RIPE or APNIC allocated networks but only import them. For more details, refer to the section [Importing SPX Allocated Networks](#).

Once you imported an allocated network, you can then add, you can add SPX assigned networks using dedicated classes, the ones that come with the appliance and/or some that you created. Note that you can also import assigned networks. For more details, refer to the section [Importing SPX Assigned Networks](#).

Keep in mind that you must indicate the SPX user that manages your network through SOLIDserver. Therefore, **before creating an assigned network, you must have a user in charge of managing it in the RIPE or APNIC database**. If the person managing the assigned network already exists in the RIPE or APNIC database, there is no need to create it in the GUI, you can import it. For more details, refer to the chapter [Managing SPX Persons](#).

Once you added an assigned network via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the assigned network addition was confirmed, you can display the column **Waiting state**. For more details, refer to the section [Browsing SPX Networks](#).

To add a SPX assigned network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **V4**.
3. In the menu, click on **+ Add an IPv4 network (subnet) by search**. The wizard opens.
4. On the page **Space selection**, select the space of your choice. Click on **NEXT**. The page **Network class** opens.
5. In the list **Network class**, select the SPX Assigned network class of your choice. Click on **NEXT**. The page **Network Size** opens.
6. Select a **Size**, **Prefix** or **Netmask**. The two other fields are edited accordingly.
7. Click on **NEXT**. The page **Search result** opens.
8. In the list **Network address**, select a start address.
9. Click on **NEXT**. The page **Add an IPv4 network** opens.
10. Configure the assigned network:

- a. The **Address** and **Prefix** fields are displayed in read-only as they correspond to the criteria previously set.
 - b. In the section **Terminal network**, the box is ticked.
 - c. In the field **Gateway**, the gateway is displayed. Its IP address corresponds to the default gateway offset configured. You can edit it if need be.
 - d. In the drop-down list **Number of pools**, you can select a value between 1 and 5, depending on the number of pools you want to create in the assigned network. Once you selected a value, you need to set the **Size** and **Type** of each pool.
 - e. In the drop-down list **Advanced properties**, *Default* is selected by default. If you want to set particular behaviors for the assigned network, select *All*. New fields appear. For more details, refer to the section [Configuring IPAM Advanced Properties](#) in the chapter [Managing Advanced Properties](#).
 - f. At the bottom of the wizard, in the field **Inetnum**, the assigned network start and end address are displayed in read-only.
 - g. In the field **Net name**, name the assigned network. The field automatically displays capital letters. The value entered in also displayed in the field *Network name*.
 - h. In the field **Description**, type in a description for the assigned network.
 - i. In the drop-down list **Country**, you can select the country where the organization is located.
11. Click on **NEXT**. The next page opens and allows you to set up a notify mail:
- a. In the field **Notify mail**, you can type in the email address of the person notified of any change made on the assigned network you are creating.
 - b. Click on **ADD**. The address is moved to the list **Notify**.
 - c. In the field **Remarks**, you can type in a comment regarding the assigned network.
12. Click on **NEXT**. The page **Contacts** opens.
- a. Specify the assigned network technical contacts (tech-c):
 1. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 2. Click on **SEARCH** to retrieve their details.
 3. Click on **ADD**. The contact is moved to the field **Technical contacts**.
 - b. Specify the assigned network administrative contacts (admin-c):
 1. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 2. Click on **SEARCH** to retrieve their details.
 3. Click on **ADD**. The contact is moved to the field **Administrative contacts**.
13. Click on **OK** to complete the operation. The report opens and closes. The assigned network is listed, its state is  *Creating*. Until its status is not  *OK*, the RIPE or APNIC has not

confirmed the addition. Have a look in the column **Waiting state** for more details regarding the addition confirmation.

If the assigned network status stays in *wait_mail_add*, refer to the section [Registering SPX Network Changes](#).

If the assigned network status stays in *wait_aw_confirm*, refer to the section [Validating a New Assignment Window](#).

To add a SPX IPv6 assigned network

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v6]**.
3. In the menu, click on **+ Add an IPv6 network (subnet) by search**. The wizard opens.
4. In the list **Choose a space**, select the space of your choice. Click on **[NEXT]**. The page **Network class** opens.
5. In the list **Network class**, select the SPX Assigned network class of your choice. Click on **[NEXT]**. The page **Network Size** opens.
6. In the drop-down list **Prefix**, select the value of your choice.
7. Click on **[NEXT]**. The page **Search result** opens.
8. In the drop-down list **Network address (v6)**, select a start address.
9. Click on **[NEXT]**. The page **Add an IPv6 network** page opens.
10. Configure the IPv6 assigned network:
 - a. The **Address** and **Prefix** fields are displayed in read-only as they correspond to the criteria previously set.
 - b. In the section **Terminal network**, the box is ticked.
 - c. In the field **Gateway**, the gateway is displayed. Its IP address corresponds to the default gateway offset configured. You can edit it if need be.
 - d. In the drop-down list **Number of pools**, you can select a value between 1 and 5, depending on the number of pools you want to create in the assigned network. Once you selected a value, you need to set the **Size** and **Type** of each pool.
 - e. In the drop-down list **Advanced properties**, *Default* is selected by default. If you want to set particular behaviors for the IPv6 assigned network, select *All*. New fields appear. For more details, refer to the section [Configuring IPAM Advanced Properties](#) in the chapter [Managing Advanced Properties](#).
 - f. At the bottom of the wizard, in the field **Inetnum**, the IPv6 assigned network start address and prefix are displayed.
 - g. In the field **Net name**, name the assigned network. The field automatically displays capital letters. The value entered in also displayed in the field **Network name**.
 - h. In the field **Description**, type in a description for the assigned network.
 - i. In the drop-down list **Country**, you can select the country where the organization is located.
11. Click on **[NEXT]**. The next page opens and allows you to set up a notify mail:

- a. In the field **Notify mail**, you can type in the email address of the person notified of any change made on the IPv6 assigned network you are creating.
 - b. Click on . The address is moved to the list **Notify**.
 - c. In the field **Remarks**, you can type in a comment regarding the assigned network.
12. Click on . The page **Contacts** opens.
- a. Specify the assigned network technical contacts (tech-c):
 1. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 2. Click on  to retrieve their details.
 3. Click on . The contact is moved to the field **Technical contacts**.
 - b. Specify the assigned network administrative contacts (admin-c):
 1. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 2. Click on  to retrieve their details.
 3. Click on . The contact is moved to the field **Administrative contacts**.
13. Click on  to complete the operation. The report opens and closes. The assigned network is listed, its state is  *Creating*. Until its status is not  *OK*, the RIPE or APNIC has not confirmed the addition. Have a look in the column **Waiting state** for more details regarding the addition confirmation.

If the IPv6 assigned network status stays in *wait_mail_add*, refer to the section [Registering SPX Network Changes](#).

Editing SPX Networks

Allocated networks cannot be edited, however, assigned networks can be edited both in IPv4 and IPv6. Any changes are sent to the RIPE or APNIC and the notify mail person configured for the assigned network.

Editing the content of your assigned networks follows the same procedures as regular assigned networks. For more details, refer to the chapters [Managing Pools](#) and [Managing IP Addresses](#).

Once you edited an assigned network via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the assigned network edition was confirmed, you can display the column **Waiting state**. For more details, refer to the sections [Browsing SPX Networks](#) and .

To edit a SPX assigned network from its properties page

1. In the sidebar, go to  **IPAM** > **Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on  or  depending on your needs. The page refreshes and the button turns black.
3. At the end of the line of the assigned network of your choice, click on . The assigned network properties pages opens.

4. In the panel **Main properties**, click on **[EDIT]**. The wizard opens.
5. In the list **Assigned network class**, select a different class or *None*.
6. Click on **[NEXT]**. The page **Edit an IPv4 Network** or **Edit an IPv6 Network** opens.
7. Edit the field **Net name**, **Description** and/or **Country**, according to your needs.
8. Click on **[NEXT]**. The next page open.
9. Edit the list of notification email addresses and Remarks field according to your needs:
 - a. Add a new email address if need be. In the field **Notify mail**, type in the new email address. Click on **[+]** to move the address in the list **Notify**. In the field **Remarks**, you can type in a comment regarding the assigned network to fill the RIPE or APNIC field *remarks*.
 - b. Remove an address from the list **Notify**. Select the address you want to delete and click on **[-]**. The address is no longer listed.
 - c. In the field **Remarks**, you can edit the comment regarding the assigned network.
10. Click on **[NEXT]**. The page **Contacts** opens.
11. You can edit the lists **Technical contacts** and **Administrative contacts** of assigned network: find new persons in the field **Nic handle / Person** and **[ADD]** them. Or select a person in the list and **[DELETE]** it.
12. Click on **[OK]** to complete the operation. The report opens and closes. The changes are listed in the panel.
13. Go to the page **All network** to see the assigned network state and make sure it was confirmed by the RIPE or APNIC. Until its status is not **OK**, the RIPE or APNIC has not confirmed the edition. Have a look in the column **Waiting state** for more details regarding the edition confirmation.

If the assigned network status stays in *wait_mail_add*, refer to the section [Registering SPX Network Changes](#).

If the assigned network status stays in *wait_aw_confirm*, refer to the section [Validating a New Assignment Window](#).

Registering SPX Network Changes

If the assigned networks you added, edited or deleted have the status *wait_mail_add*, *wait_mail_del* or *must_send_mail*, you need to use the option **Register again**. This option resends your assigned network information to the RIPE or APNIC via POST or EMAIL, depending on your configuration.

You can also send any change made on SPX persons to the RIP or APNIC. For more details, refer to the section [Registering SPX Person Changes](#).

To register again newly added or edited SPX assigned networks

1. In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **[v4]** or on **[v6]**.
3. Tick the assigned network(s) that have the status *wait_mail_add*, *wait_mail_del* or *must_send_mail*.
4. In the menu, select **[Edit] > SPX > Register again**. The wizard **Assigned network Register again** opens.

- Click on **OK** to complete the operation. The report opens and closes. The assigned network Status evolves until it is **OK**. Have a look in the column **Waiting state** for more details regarding the RIPE or APNIC confirmation.

Validating a New Assignment Window

When you add or edit an IPv4 assigned network through SOLIDserver, you can purposely exceed the Assignment Window declared during your RIPE or APNIC configuration. By exceeding, we mean:

- Configuring an assigned network which starts and/or end address exceeds the range of IP addresses available in the allocated network.
- Allocating an assigned network to a user even if this allocation exceeds the total number of IP addresses you are allowed to allocate. This sum takes into account the total number of IP addresses in your Assignment Window over the last 12 months. For more details, refer to the prerogatives in the section 7.0 Assignment Window in the document RIPE-599, available at <http://www.ripe.net/ripe/docs/ripe-599#Assignment-Window>.

In both cases, the assigned networks are marked *wait_aw_confirm*. Keep in mind that if you do exceed the AW, you need to:

- Follow the appropriate RIPE or APNIC procedures to be able to extend your Assignment Window.
- Once your request is approved by the RIPE or APNIC, you can use the option **Validate AW** in the GUI.

If your request is denied, you should delete the assigned network. For more details, refer to the section [Deleting SPX Allocated Networks](#).

To confirm the new AW in the GUI

- In the sidebar, go to **IPAM > Networks**. The page **All networks** opens.
- On the right-end side of the menu, click on **v4**.
- Tick the assigned network(s) marked *wait_aw_confirm* and approved by the RIPE or APNIC.
- In the menu, select **Edit > SPX > Validate AW**. The wizard **Assignment Window validation** opens.
- Click on **OK** to complete the operation. The report opens and closes. The page *All networks* is visible again. Have a look in the columns **Waiting state** and **Status** to monitor the evolution.

Deleting SPX Networks

From the page *All networks*, you can delete IPv4 and IPv6 SPX networks from SOLIDserver. However, note that:

- Deleting an SPX allocated network from SOLIDserver does not delete it from your RIPE or APNIC database.** It only removes it from the GUI.
- Once you deleted an assigned network via the GUI, you have to wait for the RIPE or APNIC confirmation. To make sure the assigned network deletion was confirmed, you can display the column **Waiting state**. For more details, refer to the section [Browsing SPX Networks](#).

To stop managing an SPX allocated network via SOLIDserver

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the allocated network(s) you want to delete.
4. In the menu, click on **🗑 Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The selected allocated networks are no longer listed, they might be replaced by *Orphan Networks*. **This deletion does not delete the allocated network from your RIPE or APNIC database.**

To delete an SPX assigned network

1. In the sidebar, go to **♣ IPAM > Networks**. The page **All networks** opens.
2. On the right-end side of the menu, click on **v4** or **v6** depending on your needs. The page refreshes and the button turns black.
3. Tick the assigned network(s) you want to delete.
4. In the menu, click on **🗑 Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The assigned network state is **🔄 Deleting** until the RIPE or APNIC confirms its deletion. Have a look in the column **Waiting state** for more details regarding the deletion confirmation.

If you had used addresses within a deleted assigned network, they are placed in an Orphan address and listed among your assigned networks. They are simply displayed in the GUI but no longer used within your RIPE or APNIC database as the whole assigned network was deleted.

If the assigned network status stays in *wait_mail_del*, refer to the section [Registering SPX Network Changes](#).

Chapter 86. Managing SPX AS Numbers

On the page AS Numbers of the module SPX, you can add or import the RIPE or APNIC registration details (aut-nums) for the Autonomous System (AS) Numbers that suit your needs. When it is done, all the routing routing policies, described by enumerating all neighboring AS Number with which routing information is exchanged, are listed in the page All policies. For each neighbor, the routing policy is described in terms of exactly what is being sent (announced) and allowed (accepted). That way, each aut-num contains policies that describes what can be implemented and enforced locally by said AS Number.

Note that, AS Numbers:

- Are necessary to implement BGP Anycast routing in the module DNS. For more details refer to the section [Implementing Anycast Using BGP](#).
- Can also be displayed, when used in a Route Distinguisher, on the page All routes in the module NetChange. For more details refer to the chapter [Managing Routes](#).

Browsing SPX AS Numbers

AS Numbers are visible in the module SPX, on the page AS Numbers. Any addition sends a request to the RIPE or APNIC that is confirmed or denied. The status of that request can be displayed in a dedicated column on the page.

Browsing the SPX AS Numbers Database

To display the list of SPX AS Numbers

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. On the right-end side of the menu, click on  **Listing templates**. The window opens.
3. To display the SPX AS Numbers related columns on the page, make sure the parameters described in the table below are ticked:

Table 86.1. SPX AS Numbers class parameters

Parameter	Column name
AutNum name	AutNum name
AutNum status	Status
Class param: Admin c	Administrative contacts
Class param: AS name	AS name
Class param: Description	Description
Class param: Maintainer	Maintainer
Class param: Tech c	Technical contacts
AutNum class name	Class

Note that most of these parameters are only available if you correctly enabled the classes located in the directory SPX. For more details, refer to the chapter [Enabling the SPX Classes](#).

4. Click on . The page refreshes.
5. You can filter the list to display the AS Numbers that suit your needs using the search engine of the column **AutNum name** and **AS name**.

Adding SPX AS Numbers

From the page *All AS Numbers*, you can add as many AS Numbers (aut-nums) as you need. This addition is also notified to the RIPE or APNIC.

To add a SPX AS Numbers

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. In the menu, click on **+ Add**. The wizard opens.
3. In the list **Autnum class**, select the SPX class of your choice.
4. If you or your administrator created classes at the autnum level, in the list **Autnum class** select a class or *None*. Click on **NEXT**. The page **Add an AS Number** of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Configure the AS Number:
 - a. In the field **AutNum name**, the AS Number full name is displayed once you filled in the field **AS Number** as follows: *AS<AS-number>*.
 - b. In the field **AS Number**, type in the number of your choice. This number must be available, composed of 10 digits at the most and lower than 4294967295. The value entered automatically completes the field **AutNum name**.
 - c. In the field **AS name**, you can name the AS Number.
 - d. In the field **Description**, you can type in a description.
 - e. In the drop-down list **Maintainer**, select your maintainer.
6. Click on **NEXT**. The page **Contacts** opens.
 - a. Specify the AS Number technical contacts (tech-c):
 - i. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 - ii. Click on **SEARCH** to retrieve their details.
 - iii. Click on **ADD**. The contact is moved to the field **Technical contacts**.
 - b. Specify the AS Number administrative contacts (admin-c):
 - i. In the field **Nic handle / Person**, type in the user's Nic handle or name (as displayed in the RIPE or APNIC *person* field).
 - ii. Click on **SEARCH** to retrieve their details.
 - iii. Click on **ADD**. The contact is moved to the field **Administrative contacts**.
7. Click on **OK** to complete the operation. The report opens and closes. The user is listed, its state is  *Creating*. Until its status is not  *OK*, the RIPE or APNIC has not confirmed the addition. Have a look in the column **Waiting state** for more details regarding the addition confirmation.

Editing SPX AS Numbers

There are two ways of editing an AS Number by:

1. Editing its details, i.e. AS name, Description, Maintainer and Contact information but you cannot edit the Autnum full name.
2. Editing its content, i.e. deleting some of its policies.

To edit a SPX AS Numbers

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. At the end of the line of the AS Number of your choice, click on . The properties pages opens.
3. In the panel **Main properties**, click on **EDIT**. The wizard opens.
4. If you or your administrator created classes at the autnum level, in the list **Autnum class** select a class or *None*. Click on **NEXT**. The page **Add an AS Number** of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. Click on **NEXT**. The page **Edit an AS Number** opens.
6. Edit the AS Number configuration via the fields **AS name**, **Description** and **Maintainer**,
7. Click on **NEXT**. The page **Contacts** opens.
8. Edit the contact details according to your needs.
9. Click on **OK** to complete the operation. The report opens and closes. The changes are listed in the panel.

To edit the policies of a SPX AS Numbers

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. In the column **AutNum name**, click on the name of the aut-num of your choice. The page **All policies** opens.
3. Tick the policy(ies) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on **OK** to complete the operation. The report opens and closes. The policies are no longer listed.

Deleting SPX AS Numbers

You can delete AS Numbers, it also deletes the policies it contains.

If you want to delete the policies of an AS Number refer to the section [Editing SPX AS Numbers](#).

To delete a SPX AS Numbers

1. In the sidebar, go to  **SPX > AS Numbers**. The page **All AS Numbers** opens.
2. Tick the user(s) you want to delete.

3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on  to complete the operation. The report opens and closes. The user state is  *Deleting* until the RIPE or APNIC confirms its deletion.

Part XVII. Rights Management

Managing users, groups and their authentications is an essential part of SOLIDserver configuration. Users can access any module if they belong to a group that has sufficient resources and rights. For that reason, configuring user access requires to:

1. Add or import users.
2. Create a group of users.
3. Configure that group with users. At group level, the users are managed as *Resources*, you must add them as group resource.
4. Configure that group with the rights that suit your needs. From the properties page of a group you can define which operations its users can perform (add IPAM networks, delete DHCP statics, move DNS zones...).
5. Configure that group with actual resources, i.e. existing objects from any module. Without objects as *Resources*, the users of the group can perform operations but cannot perform them on anything.

In addition, you can configure rules to authenticate external users.

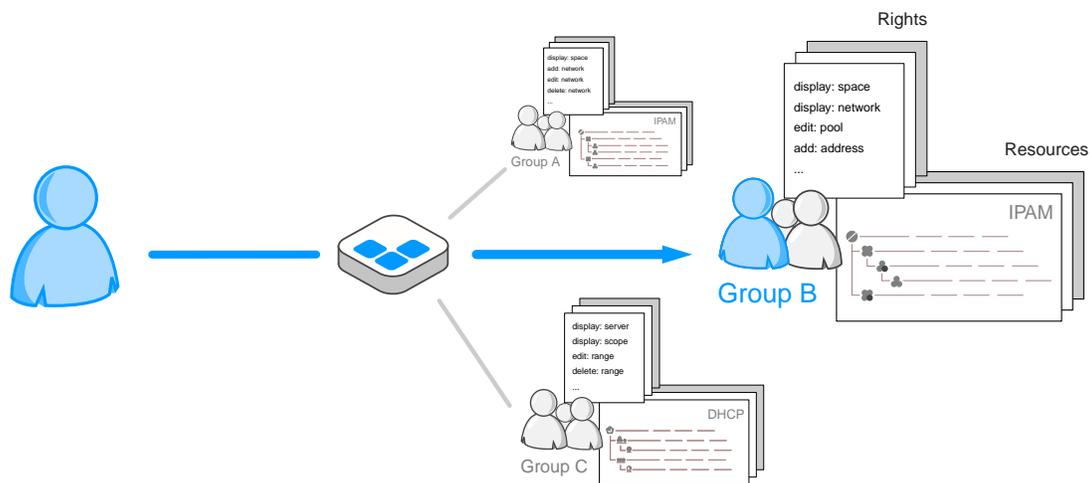


Figure 229. The rights and resources of a user depend on the group they belong to

From the module Administration, the hierarchy of Rights Management can include up to 3 levels:

- **Groups:** the highest level of the hierarchy. The groups contain resources, users and objects. For more details, refer to the chapter [Managing Groups](#).
- **Authentication Rules:** an optional second level. The rules define if external AD, RADIUS and LDAP users can access SOLIDserver. These rules allow to retrieve user credentials stored in the corresponding remote directory and to secure remote authentications. The authenticated users are granted the rights defined by the group they belong to. For more details, refer to the chapter [Managing Authentication Rules](#).
- **Users:** the lowest level of the hierarchy. You can manage local and external users on one page. Once created, you can set them as resource of a group to manage their access rights and restrictions. For more details, refer to the chapter [Managing Users](#).

Note that the superuser, *ipmadmin*, is granted all rights by default and has access to all existing objects. They belong to the most privileged group, *admin*.

Chapter 87. Managing Groups

The groups of users allow to delegate administrative rights to the users they manage. In other words, the group of users define user profiles. You can create as many groups as you want.

Indeed, the rights and resources of a group determine which operations the users can perform and upon which resources they can perform them. Any operation denied to the group is off limits to their users, any resource not listed among their resources is not available on the managing page even if the group has to access the page itself.

To successfully configure a group of users you should:

1. Create a group.
2. Grant or deny it the rights of your choice.
3. Add all the objects relevant to the granted rights as resources of the group. With rights and resources, the user profile is set.
4. Add existing users to the resources of the group. For more details, refer to the chapter [Managing Users](#).

The groups can manage remote users authenticated via RADIUS, Microsoft Active Directory or LDAP directory. For more details regarding users secure authentication, refer to the chapter [Managing Authentication Rules](#).

Browsing Groups of Users

In the rights management hierarchy, groups of users constitute the highest level. It is also the key to delegate rights.

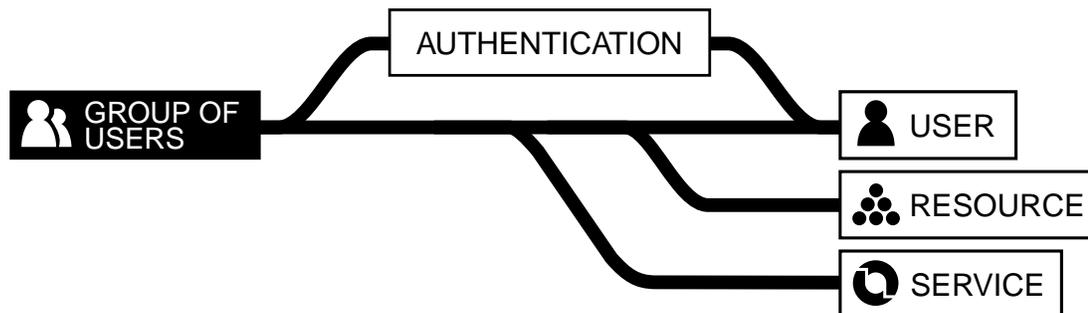


Figure 87.1. The group of users in the rights management hierarchy

By default, the group **admin** is listed on the page *Groups*. This group manages *ipmadmin*, the superuser, and is granted access to all rights and resources by default. You cannot deny it access to rights or resources. Any user added to this group has full administrative rights, like *ipmadmin*.

Browsing the Groups Database

To display the list of groups of users

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Groups**. The page opens.

To display a group of users properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Groups**. The page opens.
3. At the end of the line of the group of your choice, click on . The properties page opens.

Customizing the Display on the Page Groups

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding Groups of Users

You can add as many groups as you need. For each group you can then define users, resources and rights.

Note that you can also import groups of users from a CSV file. For more details, refer to the chapter [Importing Groups of Users](#).

We strongly suggest that you configure your group of users profiles before enabling the remote authentication rules. Once the authentication rules are enabled, the corresponding users can log in SOLIDserver. This goes especially for AD authentication: once the rule is enabled, any AD user can log in the appliance. If you created a group of users named after the AD group the users belong to, SOLIDserver automatically creates a user in the GUI and put it in the corresponding group of users. For more details, refer to the chapter [Managing Authentication Rules](#).

To add a group

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Groups**. The page opens.
3. In the menu, click on  **Add**. The wizard **Add a group** opens.
4. In the list **Parent group**, select the parent group of your choice or *None*. The selected parent group can add users to the group you are creating.
5. Click on **NEXT**. The next page opens.
6. If you or your administrator created classes, in the list **Group class** select a class or *None*. Click on **NEXT**. The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

7. In the field **Group name**, name the group. If you want users to be authenticated via AD, name the group after an existing AD group.
8. In the field **Description**, you can type in a description.
9. In the drop-down list **Copy rights from group**, you can select any group, except *admin*, or *None*. The rights of the selected group are granted to the group you are adding.

This option can be used as a template when you add new groups. If some rights should be granted or denied, you can edit the group *Rights* later on.

10. Click on **OK** to complete the operation. The report opens and closes. The group is listed.

Managing the Resources of a Group of Users

In SOLIDserver, the resources are all the objects, managed in various modules, for which you can delegate management rights. Not all objects can be set as resource.

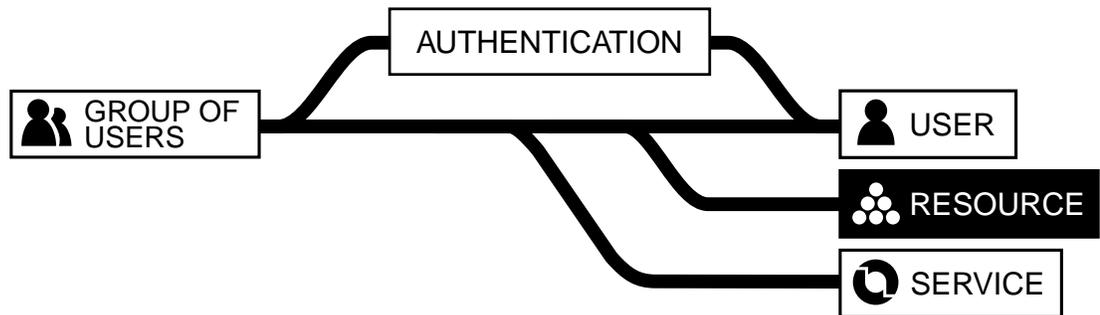


Figure 87.2. The resource in the rights management hierarchy

To display the resources of a specific group

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Groups**. The page opens.
3. Click on the **Name** of the group of your choice. The page **Resources** opens.

Following each module internal hierarchy, once an object is set as a resource the whole path in the internal hierarchy of the module is available for display. In the same way, the objects set as resource provide read-only access to lower levels.

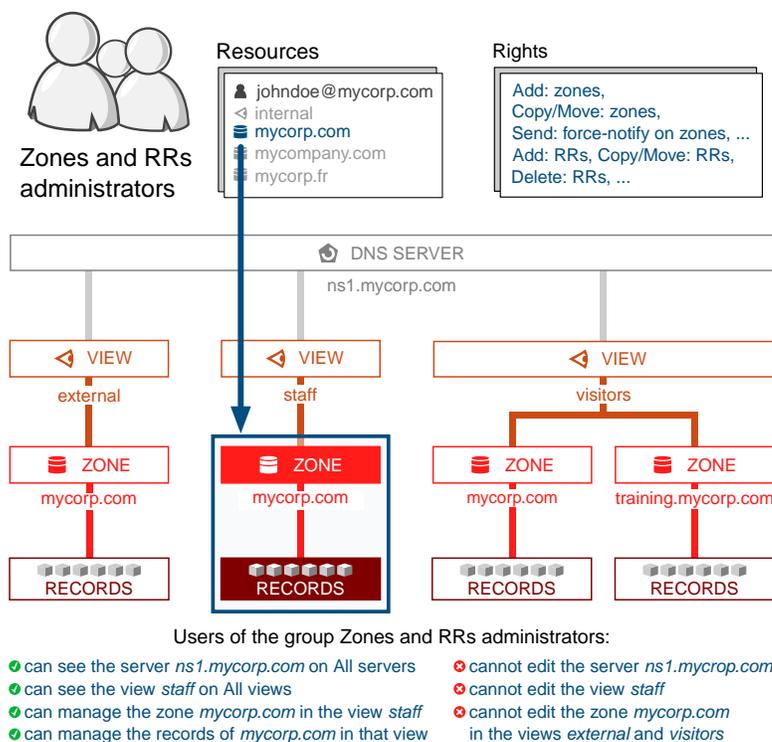


Figure 87.3. Assigning resources grants access to parts of your organization

IPAM resources

- **Spaces** as resource provide read-only access to the IPv4 and IPv6 block-type networks, IPv4 and IPv6 subnet-type networks, IPv4 and IPv6 pools and IPv4 and IPv6 addresses they contain.
- **IPv4 block-type network** as resource provide read-only access to the subnet-type networks, pools and IP addresses they contain.
- **IPv6 block-type network** as resource provide read-only access to the subnet-type networks, pools and IP addresses they contain.
- **IPv4 subnet-type network** as resource provide read-only access to the pools and IP addresses they contain.
- **IPv6 subnet-type network** as resource provide read-only access to the pools and IP addresses they contain.
- **IPv4 pool** as resource provide read-only access to the IP addresses they contain.
- **IPv6 pool** as resource provide read-only access to the IP addresses they contain.

DNS resources

- **Servers** as resource provide read-only access to the options (all) - including forwarding, access control, notify... -, views, zones, resource records, keys, access control lists and RPZ rules they contain.
- **Views** as resource provide read-only access to the options (all) - including forwarding, access control, notify... -, zones, resource records, keys, access control lists and RPZ rules they contain.
- **Zones** as resource provide read-only access to the zone options, resource records, access control lists and RPZ rules they contain.

- **RPZ zones** as resource provide read-only access to the resource records, access control lists and RPZ rules they contain.

DHCP Resources

- **Servers** as resource provide read-only access to the scopes, ranges, statics, groups, failover channels, option configurations, option definitions and ACLs they contain.
- **Servers (v6)** as resource provide read-only access to the scopes, ranges, statics, groups, failover channels, option configurations and option definitions they contain.
- **Scopes** as resource provide read-only access to the ranges, statics included in the scope and option configurations they contain.
- **Scopes (v6)** as resource provide read-only access to the ranges, statics included in the scope and option configurations they contain.

Applications Resources

- Applications, pools and nodes cannot be assigned as resources. All users have read-only access to these objects.
- **DNS servers** as resource provide read-only access to the traffic policy deployments for applications associated with the selected servers.

Guardian Resources

- Policies and triggers cannot be assigned as resources. All users have read-only access to these objects.
- **DNS servers** as resource provide read-only access to deployments of policies they are associated with.

NetChange Resources

- **Network devices** as resource provide read-only access to the ports, VLANs and discovered items they contain.

VLAN Manager Resources

- **VLAN domains** as resource provide read-only access to the VLAN ranges and VLANs they contain.
- **VLAN ranges** as resource provide read-only access to the VLANs they contain.

Administration Resources

- **Classes** as resource provide read-only access to the class objects (Class Editor) they contain.

Assigning Resources to a Group

Before assigning resources to a group keep in mind that:

- A group must be configured with resources and the rights to manage them. Without rights, at most users can see the resources but cannot perform any operations on them. For more details, refer to the section [Configuring the Rights of a Group of Users](#).
- A group managing objects configured with customized classes should have these classes among its resources. Otherwise, its users cannot see or edit the class parameters and every time they edit resources configured with this class, they delete the value of all the class parameters set for the object.
- A group of users cannot be set as resource of a group. Only users of the group *admin* can add, edit their rights or delete them.

- Users of the group *admin* have access to all resources by default. There is not need to assign it resources.
- The objects that can be set as resource provide access to their container and the objects they contain in read-only. You cannot perform operations on container or children if they are not part of your resources, they only provide an a clear overview of the objects organization.

You can add objects as resource of a group from the page *Resources*, the listing page *All <object>* or the properties page of an object.

To add resources to a group from the page **Resources**

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Click on the name of the group of your choice¹. The page **Resources** opens.
4. In the menu, select **+ Add > Resources > <resource-of-your-choice>**. The wizard opens.
5. Tick the resources you want to attribute to the group.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the selected resources are listed on the page.

On the object(s) properties page, the panel **Groups access** lists the group.

To add resources to a group from a listing page

1. From the listing page of your choice, tick the object(s) you want to set as a resource to a group.
2. In the menu, select  **Edit > Rights > Add as group(s) resource(s)**. The wizard **Resources Management** opens.
3. In the list **Available group(s)**, select a group and click on  to add the selected resources to its Resource list. The group is moved to the list **Add to the resources of the group(s)**. Repeat these actions for as many groups as needed.
4. In the list **Add to the resources of the group(s)**, the groups that have the selected objects as resource are listed. You can remove one (or several) group from that list if you do not want it to have the selected objects as a resource anymore: select the group and click on . The group is listed back in the list **Available group(s)**.
5. Click on to complete the operation. The report opens and closes. The page refreshes.

On the object properties page, the panel **Groups access** lists the group(s).

To add resources to a group from a resource properties page

1. From the listing page of your choice, display the properties page of the object of your choice using .
2. In the panel **Group access**, you can see all the groups that have the object among their resources and the actions what actions they can perform over it. Click on to add a group to the list. The wizard **Groups** opens.
3. In the field **Available group(s)**, select a group and click on  to move it to the list **Selected group(s)**. Repeat this action for as many groups as needed. All the existing groups of users

¹Any group EXCEPT the group *admin* as, by default, it has authority over all the resources of SOLIDserver database.

are listed except *admin* as all the objects of the database are resources of the group by default.

4. In the field **Selected group(s)**, the groups that have the object as a resource are listed.
5. Click on to complete the operation. The report opens and closes. The page refreshes and the panel is updated.

Removing Resources from a Group

At any time you can remove objects from a group list of resources. This prevents the users of the group from accessing them unless a container or lower object of the same hierarchy is listed among their resources.

For instance, if a group has among its resources the space *Local* and a network *local-terminal-network*, when you decide to remove access to the space, you can remove it from their resources. Keep in mind that removing the space prevents users from performing operations on the space but they can still access its content because the network *local-terminal-network* is still part of their resources: users have read-only access to the resource path within the module.

Keep in mind that to remove resources from a group, users that do not belong to the group *admin* need to be granted the appropriate rights.

To remove a resource from a group

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Click on the name of the group for which you want to remove resources. The page **Resources** opens.
4. Tick the box of the resource(s) you want to delete.
5. In the menu, click on  **Delete**. The wizard **Delete resources** opens.
6. Click on to complete the operation. The report opens and closes. The selected resources are no longer listed.

Managing the Rights of a Group of Users

The rights are services in essence that can be granted or denied to groups of users.

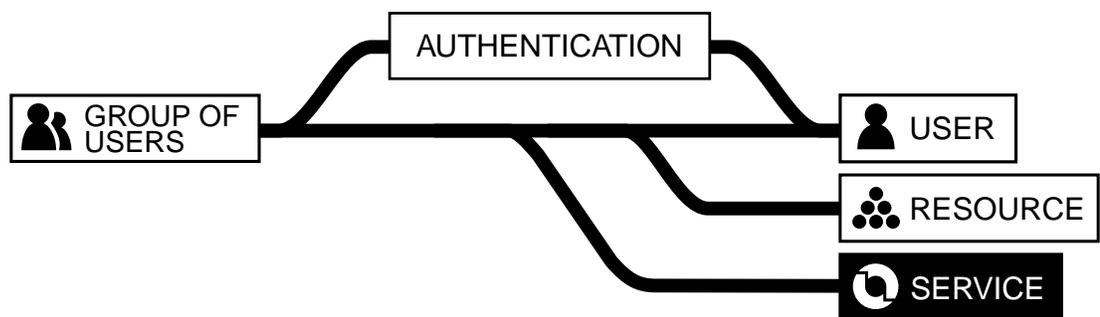


Figure 87.4. The service in the rights management hierarchy

Before configuring rights keep in mind that:

- Only users from the group *admin* can configure the rights of the other groups of users.
- The users of each group are granted or denied access to the rights configured for the group(s) they belong to.
- Rights alone are useless. Without any resource, the users of the group cannot use the rights they are granted. For more details, refer to the section [Assigning Resources to a Group](#).
- The rights are managed one group at a time from the page **Rights**. The page displays all the configurable rights and their details in each column. The rights do not have a properties page.
- The group *admin* is the only one with full administrative rights over SOLIDserver. All the users of that group, including *ipmadmin*, are granted all the rights by default and can perform all the tasks. You cannot remove any rights from this group.
- All the users have access to the page **User tracking**, where they can view all the operations they performed in every module they have access to.

You can grant a group access to the operations performed by all SOLIDserver users if you grant it access to the User tracking right.

The page **Rights** contains the following columns:

Table 87.1. The columns on the page Rights

Column	Description
Permission	The permission status of a right:  Allow or  Deny.
Right	The name of the right following the format: <operation>: <object-concerned>.
Module	The module or type of the right: <i>IPAM, DHCP, DNS, NetChange, Workflow, Device Manager, VLAN Manager, VRF, SPX, Administration, Rights & delegation</i> or <i>Reports</i> .
Access type	The type of right: <i>Read</i> or <i>Write</i> .

The name of each right contains a verb that detail the *operation*. The most used verbs are the following:

Table 87.2. The most common operations in the name of the rights

Verb	Description
Add	The service allows to add and edit an object.
Delete	The service allows to delete an object.
Display	The service allows to display the complete list of objects on its management page.
Edit	The service allows to perform specific edition operations on an object.
List	The service allows to delegate administrative rights from one administrator to the other, usually group management related rights.

There are more operations available: *Configure, Convert, Copy/Move, Display, Migrate, Split...*

Configuring the Rights of a Group of Users

From a group properties, you can set the rights to shape the users of the group profile. By default upon creation, a group of users has all read rights.

The other rights, to add, edit, delete... objects must be granted specifically in each module panel of the properties page. For instance, if you grant a group the right to edit networks but did not assign them any network, the users of the group have access to the page *All networks* and to the menu *Edit* but cannot see the list of existing networks. Hence the need to grant right AND assign resources.

To configure the rights of a group

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Click on the **Name** of the group of your choice. The page **Resources** opens.
4. In the breadcrumb, click on **Rights**. The page opens.
5. Filter the list if need be.
6. Tick the right(s) of your choice. The menu *Edit* is now selectable.
7. In the menu, select **Edit > Allow** or **Deny**. The wizard opens.
8. Click on to complete the operation. The page refreshed, the column Permission displays the current configuration of the right.

At any time, you can edit this configuration using the same procedure.

Managing the Users of a Group of Users

The purpose of a group is to define a set of rights and resources for the users it contains. Once you configured the group, you can add the existing users that fit the group profile.

Adding a User to a Group

Once you created a user, you can add it as the resources of any group. It can also belong to several groups with different resources and rights. The user credentials are the same but their access correspond to the group they belong to.

For more details regarding user creation, refer to the chapter [Managing Users](#).

To add a user to a group

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Click on the name of the group of your choice. The page **Resources** of that group opens.
4. In the menu, select **+ Add > Users**. The wizard **Rights & delegation: Users** opens.
5. Tick the user(s) you want to add to the group.
6. Click on to complete the operation. The wizard closes and the page refreshes. The user is listed among the resources of the group.

Removing a User from a Group

To restrict user rights, you can remove them from a group. If they do not belong to a group, their credentials can open SOLIDserver GUI but they cannot display or edit any module or resource.

To remove a user from a group

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.

3. Click on the name of the group of your choice. The page **Resources** of that group opens.
4. Filter the list if need be.
5. Tick the user(s) you want to remove from the group.
6. In the menu, click on  **Delete**. The wizard **Delete** opens.
7. Click on to complete the operation. The report opens and closes. The user(s) is no longer listed in the resources so can no longer benefit from the group rights. The user is still listed on the Users page.

Editing Groups of Users

Editing a group means editing the group main properties as well as the group access and restrictions. For more details regarding user rights, refer to the section [Managing the Rights of a Group of Users](#) below.

To edit a group main properties

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. At the end of the line of the group you want to edit, click on . The properties page opens.
4. In the panel **Main properties**, click on . The wizard **Edit a group** opens.
5. If you or your administrator created classes, in the list **Group class** select a class or *None*. Click on . The last page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Group name**, name the group. If you want users to be authenticated via AD, name the group after an existing AD group.
7. In the field **Description**, you can type in a description.
8. In the drop-down list **Copy rights from group**, you can select any group, except *admin*, or *None*. The rights of the selected group are granted to the group you are adding.

This option can be used as a template when you add new groups. If some rights should be granted or denied, you can edit the group *Rights* later on.

9. Click on to complete the operation. The report opens and closes. The properties page is visible again and includes the changes in the panel.

Enabling or Disabling Groups of Users

By default when you add groups, they are enabled but you can disable them.

Keep in that **if you disable a group, the users its contain are still able to connect to SOLIDserver but do not have access to any module or resource.**

To enable/disable a group

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Tick the group(s) of your choice.
4. In the menu, select  **Edit > Status > Enable** or **Disable**. The corresponding wizard opens.
5. Click on to complete the operation. The report opens and closes. The group(s) is marked *OK* or *Disabled* in the column **Status**.

Deleting Groups of Users

At any time you can delete a group.

Keep in mind that **if you delete a group, the users its contain are still able to connect to SOLIDserver but do not have access to any module or resource.**

To delete a group

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. Tick the group(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The group is no longer listed.

Chapter 88. Managing Users

The notion of user allows to define administrator accounts vs. standard user accounts. It also allows to set up different profiles and levels of management.

By default, users authentication is performed using the local database. However you can create:

Local users

If you want to use local authentication, you must configure a group and manually add this local user into the group. Once added to a group, a user is considered as a resource of the group. For more details, refer to the section [Managing the Users of a Group of Users](#) of the chapter Groups.

Remote users

If you want to authenticate users remotely via LDAP, Active Directory and RADIUS directory in addition to local authentication, you must configure those services on SOLIDserver. A local user and a remote user cannot share the same login account: if a user is already declared in the local database, it can no longer be authenticated externally. SOLIDserver comes with an authentication subsystem that manages authentications to securely log in to its web user interface. For more details, refer to the chapter [Managing Authentication Rules](#).

If you are using remote authentication, you must always have at least one local *admin* user in a local group to ensure connectivity to SOLIDserver in case your remote directory is unreachable.

If you want to manage RIPE or APNIC persons, refer to the part [SPX](#).

Note that any connected user can set their display settings and change their password, as detailed in the section [Account Configuration](#).

Browsing Users

As far as the rights management is concerned, users constitute, along with objects, the second level of the rights delegation management. Users are merely created and configured to be managed by one or several groups that set their profile and rights.

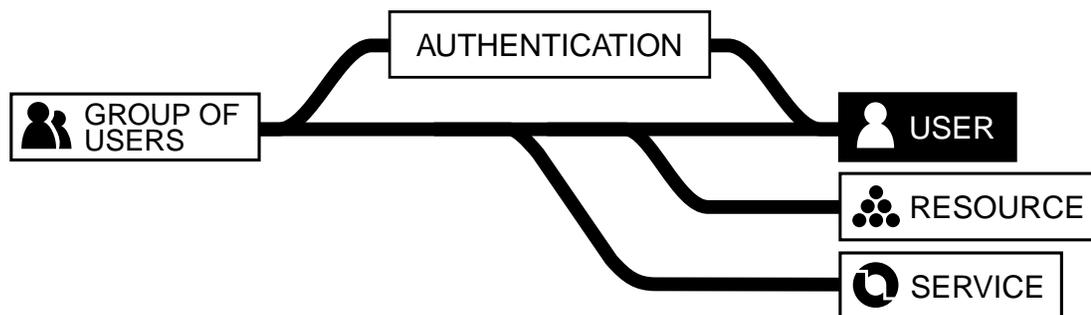


Figure 88.1. The user in the rights management hierarchy

By default, the superuser *ipmadmin* is the only user listed on the page. It belongs to the *admin* group and has all the rights and every manageable object among its resources.

Browsing the Users Database

To display the list of users

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.

To display a user properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. At the end of the line of the user of your choice, click on . The properties page opens.

Customizing the Display on the Page Users

Users of the group *admin* can create customized column layouts. The button  **Listing template**, on the right-end side of the menu, allows any user to display them. For more details, refer to the section [Customizing the List Layout](#).

Adding Users

You can add as many local users as you want on the page *Users*. Once listed on the page, the users are part of SOLIDserver local database.

Keep in mind that each local user profile depend on:

1. The group or groups, of users they belong to.
2. The right granted or denied to the group(s).
3. The resources granted or denied to the group(s).

Note that you can also import local users from a CSV file. For more details, refer to the chapter [Importing Users](#).

To add a local user

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. In the menu, click on  **Add**. The wizard **Add a user** opens.
4. If you or your administrator created classes, in the list **User class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

5. In the field **Login**, type in the user login. This login **cannot** be an email address.
6. In the field **Password**, type in the user password¹.
7. In the field **Confirm password**, type in the user password again.

¹If the user is of Unix type and the password is not printable, the system password is used.

8. If need be, configure additional parameters:
 - a. Tick the box **Expert mode**.
 - b. In the field **First name**, type in the user first name.
 - c. In the field **Last name**, type in the user last name.
 - d. In the field **Pseudonym**, the user last and first name are automatically displayed. You can replace them by a shortname or shorter name if you want.
 - e. In the field **Email**, type in the user email address.
 - f. In the field **Login URL**, type in the URL toward which the user should be directed after being authenticated.
 - g. In the drop-down list **Maintainer group**, select the group of users that should be able to edit the user information (names, credentials, email...) and classes.
9. Click on to complete the operation. The report opens and closes. The user is listed among the users with its **Login**, **Official name** and **Origin** in the corresponding columns.

Connected users can edit their session time and date or listing page display, interface language or password. For more details, refer to the section [Account Configuration](#).

Editing Users

At any time an administrator can edit a user details or group.

Keep in mind that users must belong to a group: a user not belonging to any group, they can connect to SOLIDserver but their session does not contain any module or they cannot perform any action as no right is granted to individual users.

Editing the User Details

Editing the user details means editing the *Main properties* on the user properties page.

To edit a local user information

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. In the column **Login**, click on the user name of your choice. The properties page opens.
4. In the panel **Main properties**, click on . The wizard **Edit a user** opens.
5. If you or your administrator created classes, in the list **User class**, select a class or *None*. Click on . The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. Edit the user information according to your needs. You can edit any detail except their *Pseudonym*, all parameters are detailed in the procedure [To add a local user](#).

If you type in a different password from the initial one, you overwrite the user current password. Overwriting a user password can log them out after your changes, they may no longer be able to log in.

7. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and includes the changes in the panel.

Editing the User Group

Editing the user details means modifying the panel Groups access on the user properties page. Keep in mind that in the procedure below, the changes are performed from the button *EDIT* in the panel, but you can also make these changes from the menu *Edit*.

To edit a local user group

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. In the column **Login**, click on the user name of your choice. The user properties page opens.
4. In the panel **Groups access**, click on **EDIT**. The wizard **Groups** opens.
5. If you or your administrator created classes, in the list **User class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the list **Available group(s)**, you can select a group and click on **+**. The group is moved to the **Selected group(s)**.
7. In the list **Selected group(s)** are displayed the group(s) the user belongs to. In other words, the user profiles. You can remove a group from the list clicking on **-**, the group is moved to the list **Available group(s)**.
8. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again and includes the changes in the panel.

Changing the User Password

At any time an administrator can change a user password, from the user properties page.

If you want to set a specific password complexity, refer to the section [Configuring the User Password Complexity](#).

To change a user password

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. In the column **Login**, click on the user name of your choice. The properties page opens.
4. In the panel **Main properties**, click on **EDIT**. The wizard **Edit a user** opens.
5. If you or your administrator created classes, in the list **User class**, select a class or *None*. Click on **NEXT**. The next page of the wizard opens.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **Password**, type in the new password.

3. In the column **Name**, type in *ipmserver.login.bad_login_retry_before_freeze* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value set for the key. The wizard **Registry database Edit a value** opens.
5. In the field **Name**, the key name is displayed in a read-only gray field.
6. In the field **Value**, set the maximum number of failed connection attempts, from a same IP address, after which these attempts are locked out.
7. Click on to complete the operation. The report opens and closes. The new value is displayed.

To edit the registry key to set the authentication period

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *ipmserver.login.bad_login_test_window* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value set for the key. The wizard **Registry database Edit a value** opens.
5. In the field **Name**, the key name is displayed in a read-only gray field.
6. In the field **Value**, set the authentication period during which the maximum number of failed connection attempts is allowed.
7. Click on to complete the operation. The report opens and closes. The new value is displayed.

To edit the registry key to set the lock out period

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the column **Name**, type in *ipmserver.login.bad_login_freeze_time* and hit Enter. The key is the only one listed.
4. In the column **Value**, click on the value set for the key. The wizard **Registry database Edit a value** opens.
5. In the field **Name**, the key name is displayed in a read-only gray field.
6. In the field **Value**, set the number of seconds during which connection attempts are locked out after the maximum number of failed connection attempts has been reached.
7. Click on to complete the operation. The report opens and closes. The new value is displayed.

To display the failed connections logs

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
3. In the drop-down list **SOLIDserver**, verify that the local appliance is selected. Only the host-name appears with no IP address.
4. In the drop-down list **Services**, select *auth*. The page refreshes.

5. In the column **Log**, type in *invalid* and hit Enter. The repeatedly failed attempts to connect are listed, as in the example below where the invalid user name is *toto*:

```
Feb 23 16:09:36 solid ipmsserver[1226]: invalid login/password toto
Feb 23 16:09:37 solid ipmsserver: Last message 'invalid login/passwo' repeated 2 times, suppressed
by syslog-ng on solid.intranet
Feb 23 16:09:37 solid httpd: SOLIDserver: too many connection attempts with invalid credentials.
Retry later.
```

Configuring User Sessions

SOLIDserver provides the possibility to set an automated session logout of any user if they do not do anything on the server after a certain period of time. You can also redirect users after their session expires or when they log out.

Configuring Users Login Session Time

By default, the session is set to 0, i.e. the user session does not end unless the user logs out. You cannot set the login session to less than 1 minute.

To set up a time limit for a user login session

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine column **Name**, type in *www.login.session_timeout*. Only this key is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, set the session time of your choice, in seconds. This value can be 0 or anything above 60. By default, it is set to 0.
6. Click on to complete the operation. The report opens and closes. The new value is visible in the list and now the user is automatically logged out if no actions are performed above the number of seconds you just set.

Redirecting Users After They Log Out or Their Session Expires

By default, once a user session expires or after a log out, the login page opens. You can choose to redirect users toward a website instead. A registry database entry can be configured with the website URL.

To redirect users at the end of their session

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine column **Name**, type in *logout.session.redirect.url*. Only this key is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.

5. In the field **Value**, type in the URL of your choice following the format *http://<website-of-your-choice>*. By Default, the value is *empty*.
6. Click on to complete the operation. The report opens and closes. The new value is visible in the list and now all users are automatically redirected to the website specified as value once they log out.

Enabling or Disabling Users

By default adding users, enables them. Disabling users prevents them from connecting to SOLIDserver, without deleting them from the database. You can enable them again at any time.

To enable/disable a user

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page opens.
3. Tick the user(s) of your choice.
4. In the menu, select  **Edit > Status > Enable** or **Disable**. The corresponding wizard opens.
5. Click on to complete the operation. The report opens and closes. The user(s) is marked *OK* or *Disabled* in the column **Status**.

Generating User Reports

EfficientIP provides one user dedicated report.

Table 88.1. Available user report

Page	Report
Users	Users rights in each group

For more details regarding the reports and their generation, refer to the section [Managing Reports](#).

Deleting Users

Deleting local users prevents them from connecting to SOLIDserver and removes them from the page *Users*.

Keep in mind that deleting users connecting remotely, like AD users, does not prevent them from connecting. Once the rule is enabled, users are created locally upon connection and placed in an existing group of users if their name matches the name of the group they belong to in the Active Directory.

To delete a local user

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Users**. The page **Users** opens.
3. Tick the user(s) you want to delete.
4. In the menu, click  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The user(s) is no longer listed.

Chapter 89. Managing Authentication Rules

SOLIDserver provides remote authentications management to securely log in external users to the GUI. From the page *Authentication rules*, you can set up the access for users whose credentials are stored on **Microsoft Active Directory**, **LDAP** and/or **RADIUS** servers. To set up remote authentication for SSH connections, refer to the appendix [Using Remote Authentication for SSH Connections to SOLIDserver](#).

To provide secure user connections, all remote authentications are challenged when a user connects with their login and password. If several authentications rules are configured, the connection process is the following:

1. The first authentication rule is used to authenticate the user. If it fails, SOLIDserver tries the next authentication rule. Each configured authentication rule is tried and used, whether it relies on AD, LDAP or RADIUS, until it is successful or all authentications fail.
2. Either the user is granted access, or all rules failed and the user access to the GUI is denied.
3. When the authentication succeeds, the user rights are based on the group they belong to in the remote directory. SOLIDserver tries to match the local database entries with the remote information specified in the rules. If matching names are found, they are used to allow or deny user access to operations and resources in the GUI. If no local group name matches, the connection is denied.

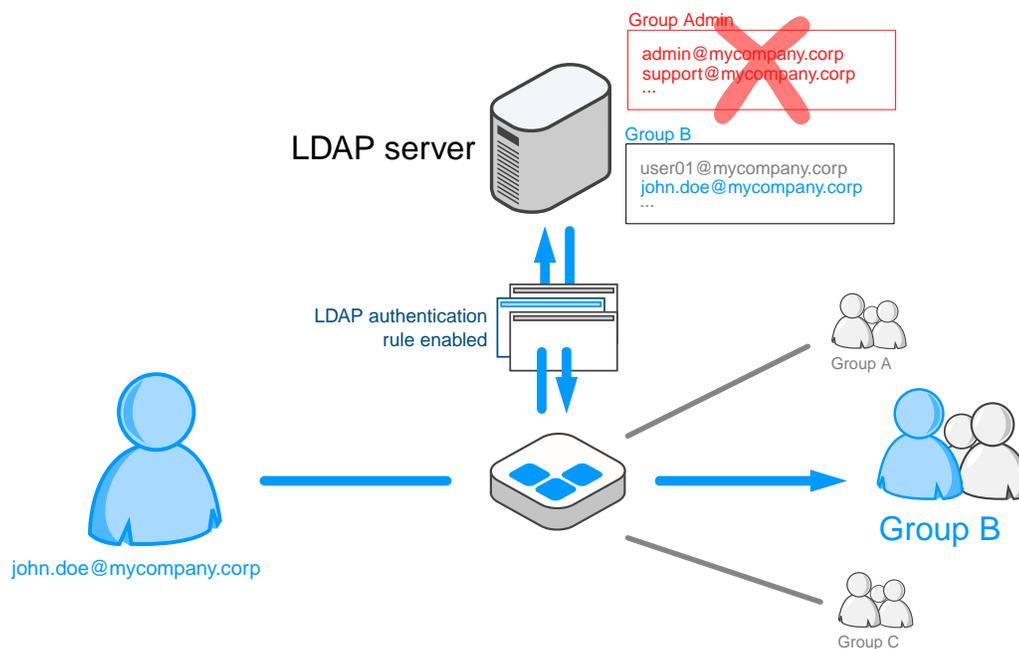


Figure 89.1. The user connection process when an authentication rule is enabled

To remedy any issue you could have when setting up remote authentication, keep in mind that you should always have at least one local admin user configured. Having a user in the group *admin* allows to have access to SOLIDserver and intervene in case of authentication failures,

if an unreachable remote directory or server, or even if a local group of users matching a remote group has insufficient rights or resources.

You can add remote authentication rules, delete, disable or enable them again from the page *Authentication rules*. Only the rules you add are displayed on the page. Keep in mind that these rules are also listed to the page *Rules* that lists all SOLIDserver rules.

Browsing Authentication Rules

As far as rights management is concerned, the authentication rules are optional and to be used only if you intend to allow authentication of remote users through AD, LDAP or RADIUS.

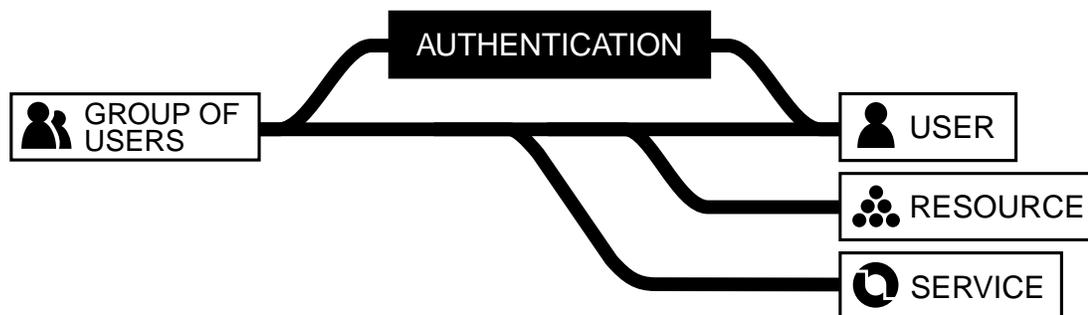


Figure 89.2. The authentication rule in the rights management hierarchy

By default, the list is empty.

Browsing the Authentication Rules Database

To display the list of authentication rules

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.

To display an authentication rule properties page

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. At the end of the line of the rule of your choice, click on **ⓘ**. The properties page opens.

Adding Authentication Rules

From the page *Authentication rules* you can add three rules dedicated to remote users authentication whether your authentication relies on Active Directory, LDAP or RADIUS. As we saw in the introduction, if your users credentials are saved on several or all of these remote servers, you can add several or all rules as SOLIDserver compare the user credentials to any identified remote server configure when adding the rules to provide the secure authentication of the user or deny it if the user is not found anywhere.

Keep in mind that thanks to this systematic check of all the remote authentication possible configurations, you can add as many rules as you want. They are all checked against the user credentials. This allows you to set different configurations for LDAP, RADIUS or AD authentication of the remote users.

Note that there is no order in which authentication rules are checked, therefore, it is important to keep your authentication servers updated with the same data. If a user is in the group A on a server and in the group B on another server, there is no mean to set a preference for one authentication rule over the other.

Relying on Active Directory Authentication

You can add the rule *AD authentication* that configures the users authentication through a Microsoft Active Directory server. Active Directory (AD) is a technology created by Microsoft that provides a variety of network services, including LDAP like directory services and other network information. SOLIDserver supports remote authentication with any AD running on Microsoft Window Server 2008, 2008 R2, 2012 R2 or 2016.

In order to use the AD authentication successfully, the following **prerequisites** must be met:

1. At least one group exists both on the AD server and in SOLIDserver database. They must have exactly the same name and this name is case sensitive: so the name of the group in SOLIDserver must respect the AD group name.
2. The user you use for testing the authentication has to be part of the group mentioned above.

To add the AD users authentication rule

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add a rule** opens.

The field **Module** is already filled with *Rights & delegation*. The field **Event** is already filled with *External user login*.

4. In the drop-down list **Rule**, select *(000) AD authentication*.
5. In the field **Rule name**, name the rule.
6. In the field **Comment**, you can add a comment regarding that rule.
7. Click on **[NEXT]**. The page **Rule filters** opens.
8. Click on **[NEXT]**. The page **Rule parameters** opens. To only configure basic authentication, refer to the step 10. If you want to configure the authentication in more details, refer to the steps 10 and 11.
9. Configure the basic AD authentication parameters following the table below:

Table 89.1. Active Directory basic parameters

Field	Description
AD server	The IP address or hostname of the Active Directory server.
AD server port	The port of the Active Directory server. Leave it empty to use the default AD port.
Use secure LDAP	Tick this box to use secure LDAP during the authentication challenge. SOLIDserver uses LDAP and SSL to connect to the AD server ^a .

Field	Description
Domain of DC	The domain of the Domain Controller. Fully Qualified Domain Name (FQDN) of your AD, for instance <i>mydomain.corp</i>
Default user domain	The default domain of the user who connects through AD. This domain is concatenated to the user name. For instance, the user login <i>jdoe</i> is concatenated with <i>mydomain.corp</i> to produce <i>jdoe@mydomain.corp</i> . If you let this field empty, you have to connect with <i>jdoe@mydomain.corp</i> . If you configure <i>mydomain.corp</i> in this field then you only have to connect with <i>jdoe</i> .

^aSecure LDAP configuration relies on the TCP port 638. Make sure this port is open on your network.

10. Configure the advanced AD authentication parameters following the steps below:
 - a. Configure the basic parameters following the [Active Directory basic parameters](#) above.
 - b. Tick the box **Expert mode**. The remaining configuration fields appear.
 - c. Finish the configuration following the table below:

Table 89.2. Active Directory expert mode parameters

Field	Description
Deny if not in a group	Select <i>Yes</i> if you only want members of an AD group to be able to connect to SOLIDserver. By default, <i>No</i> is selected. This field is optional.
Manage imbricated groups	Tick this box if to allow SOLIDserver to look for members in sub-groups of the specified top group on the AD server during the authentication challenge.

11. Click on **NEXT**. The last page of the wizard opens.
12. In the drop-down list **Synchronize**, you can choose to synchronize or not SOLIDserver database with the AD database: this automatically puts users in the local group of users that matches the AD group name they belong to. This grant them the rights of said local group.

If you select *Yes*, the box **Expert mode** appears. You can tick it to configure specific synchronization parameters, described in the table below:

Table 89.3. Active Directory parameters for the groups synchronization

Field	Description
AD group granted "admin" rights	Type in the name of any group on the AD server. All the users of the specified group are granted access to SOLIDserver with the same rights as the users of the group <i>admin</i> . These users are also listed as resource of the group <i>admin</i> .
Login	Specify an account to be used to browse AD attributes. If your AD is configured in a very strict manner, standard users might not be able to browse their own attributes. Filling this field enables SOLIDserver to retrieve the groups the user belongs to using the account specified. This field is optional, as the AD user accounts might have the sufficient privileges to browse their own attributes and retrieve the groups it belongs to.
Password	If you specified an account in the field <i>Login</i> above, type in its account password. This field is compulsory if you specified a login.
Base DN	Type in the name of the top of the AD tree. The level specified is the starting point of the search for a matching user account on the server. You can customize this field in order to look in specific location(s) of the AD. This field is optional.
Use sAMAccountName field as login	You can decide to use or not the sAMAccountName field as user login. This parameter is used for pre-AD installation (basically NTDS) and accepts 8-characters long login names instead of regular longer names. This field is optional.

13. Click on **OK** to complete the operation. The report opens and closes. The rule is listed. In the column **Instance**, the *Rule name* you chose is displayed.

Once the rule is added, AD users can connect to SOLIDserver. This connection automatically creates the user and puts in the corresponding group if you chose to synchronize the groups.

Also, note that:

- The changes performed on the AD server are not immediately taken into account by SOLIDserver. To avoid waiting, you can delete the AD users you modified from the page **Users**, when they connect again, SOLIDserver contacts the AD server and authenticates them with their new parameters.
- If several email addresses are available for a same user, only the first non-empty value is taken into account.

If some users connections fail, here below are some useful guidelines to follow.

How to troubleshoot a remote AD authentication

1. **Log out of the system** then try to logging in again. It should work, if not:
2. **Check the Syslog page** and look for any AD related information. Most of the time, the problem is coming from:
 - a. The AD connection is not possible: you should see messages telling you the `ldap_bind` was not possible.
 - b. The AD user credentials are not recognized as a member of any group SOLIDserver knows.

Relying on LDAP Authentication

You can add the rule *LDAP authentication* that configures the users authentication through LDAP (version 2 or 3). Lightweight Directory Access Protocol (LDAP) is an application protocol over TCP/IP for querying and modifying directory services that might hold passwords, addresses, groups, public encryption keys and other exchange-facilitating data.

To add the LDAP users authentication rule

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. In the menu, click on **+ Add**. The wizard **Add a rule** opens.

The field **Module** is already filled with *Rights & delegation*. The field **Event** is already filled with *External user login*.

4. From the list **Rule**, select *(018) LDAP authentication*.
5. In the field **Rule name**, name the rule.
6. In the field **Comment**, you can add a comment regarding that rule.
7. Click on **NEXT**. The page **Rule filters** opens.
8. Click on **NEXT**. The page **Rule parameters** opens.

9. Configure the rule parameters following the table below:

Table 89.4. LDAP parameters

Field	Description
LDAP server	Type in the IP address or hostname of the LDAP server. This field is mandatory.
LDAP server port	The port of the LDAP server. Leave it empty to use the default LDAP port.
Use secure LDAP	Tick this box to use secure LDAP during the authentication challenge. SOLIDserver uses LDAP and SSL to connect to the LDAP directory.
Use LDAP v3	Tick the box to use LDAP in version 3 ^a .
Base DN	Type in the top level of the LDAP directory tree is the base, as follows: <i>dc=example,dc=com</i> . This field is mandatory.
Group attribute	Type in the name of the attribute in LDAP that matches one or several groups in SOLIDserver, for instance <i>memberof</i> . The names must be separated by a comma. This field is optional.
LDAP group granted "admin" rights	Type in the name of any group on the LDAP server. All the users of the specified group are granted access to SOLIDserver with the same rights as the users of the group <i>admin</i> . These users are also listed as resource of the group <i>admin</i> . This field is optional.
Login	Type in the login ^b of an account that has sufficient privileges to retrieve user attributes during the authentication. If your LDAP standard users cannot browse their attributes, they should not be able to connect to SOLIDserver on their own. This field is optional.
Password	Type in the password of the account specified in the field <i>Login</i> above. This field is optional.

^aNot ticking this box means using LDAP in version 2.

^bIt is based on the LDAP attribute *uid*.

10. Click on to complete the operation. The report opens and closes. The rule is now listed. In the column **Instance**, the *Rule name* you chose is displayed.

Once the rule is added, LDAP users can connect to SOLIDserver. This connection automatically creates the user and puts in the corresponding group if you chose to synchronize the groups.

Also, note that:

- The changes performed on the LDAP server are not immediately taken into account by SOLIDserver. To avoid waiting, you can delete the LDAP users you modified from the page Users, when they connect again, SOLIDserver contacts the LDAP server and authenticates them with their new parameters.
- If several email addresses are available for a same user, only the first non-empty value is taken into account.

Relying on RADIUS Authentication

You can add the rule *RADIUS authentication* to configure the users authentication through any RADIUS server. Remote Authentication Dial In User Service (RADIUS) is a networking protocol that uses access servers to provide centralized access management to large networks.

Note that, when you using RADIUS authentication:

- A user with no access to RADIUS cannot access SOLIDserver either.

- The group name sent by RADIUS has to be exactly the same as the one configured in SOLIDserver, including the letter case and potential accents. RADIUS return value can hold multiple values, i.e. several groups, separated by a comma.
- RADIUS users are automatically added to SOLIDserver when connecting for the first time.
- Users attributes, such as group or email, are updated at each connection.

You can use FreeRADIUS or RADIUS for Cisco ACS with SOLIDserver. For more details refer to the appendix [Configuring RADIUS](#).

To add the RADIUS users authentication rule

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. In the menu, click on **Add**. The wizard **Add a rule** opens.

The field **Module** is already filled with *Rights & delegation*. The field **Event** is already filled with *External user login*.

4. In the list **Rule**, select *(017) RADIUS authentication*.
5. In the field **Rule name**, name the rule.
6. In the field **Comment**, you can add a comment regarding that rule.
7. Click on **NEXT**. The page **Rule filters**.
8. Click on **NEXT**. The page **Rule Parameters**.
9. Configure the rule parameters following the table below:

Table 89.5. RADIUS authentication rule parameters

Field	Description
RADIUS server IP address	Type in the IPv4 address of the host server.
RADIUS server port	Type in the port number of the UDP port used to contact the RADIUS server. If you type in the port 0, the library looks up the radius/udp or the radacct/udp service in the network services database and use the port found there. By default, the RADIUS server port used for authentication is 1812.
RADIUS secret passphrase	Type in your RADIUS password. This password is necessary to grant SOLIDserver access to RADIUS.
RADIUS request timeout (seconds)	Set up the timeout parameters. In other words, choose after how many seconds you want your radius server to switch to timeout status if no reply is received past this period of time. By default, the number of seconds is 3.
RADIUS max tries before giving up	Set the maximum number of requests to be sent before the server stops trying to connect and switches to failure state. By default, the number of retries is 3.
RADIUS NAS IP address	Type in the IP address that SOLIDserver needs to connect to RADIUS.

10. Click on **OK** to complete the operation. The report opens and closes. The rule is now listed. In the column **Instance**, the *Rule name* you chose is displayed.

Editing Authentication Rules

You might need to change the rules parameters. The procedure below details the process from the page *Authentication rules* but you can also do it from the rule properties page Main properties panel.

To edit a user authentication rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. Filter the list if need be.
4. Right-click over the **Name** or **Instance** of the rule you want to edit. The contextual menu opens.
5. Click on . The wizard **Edit a rule** opens.
6. Edit the **Rule name**, **Comment** fields and any other fields and configurations according to your needs. For more details, refer to each authentication rule addition procedure in the [Adding Authentication Rules](#) section above.
7. Click on  to complete the operation. The report opens and closes. The rule is listed. In the column **Instance**, the *Rule name* you chose is displayed.

Enabling or Disabling Authentication Rules

Once added, the authentication rules are automatically enabled. You can disable and enable them back as you please.

To enable/disable a user authentication rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. Tick the rule of your choice.
4. In the menu, select  **Edit > Enable** or **Disable**. The corresponding wizard opens.
5. Click on  to complete the operation. The report opens and closes. The rule is listed and marked  *OK* or  *Disabled* in the column **Status**.

Deleting Authentication Rules

If you no longer need an authentication rule, you can delete it.

To delete a user authentication rule

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Authentication rules**. The page **Authentication rules** opens.
3. Tick the rule of your choice.

4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on  to complete the operation. The report opens and closes. The rule is no longer listed.

Part XVIII. Administration

The module Administration should be handled by an administrator as it allows to manage remote appliances and monitor, maintain or upgrade SOLIDserver. Its homepage, *Admin Home*, is divided into six sections that contain links toward all the pages of the module.

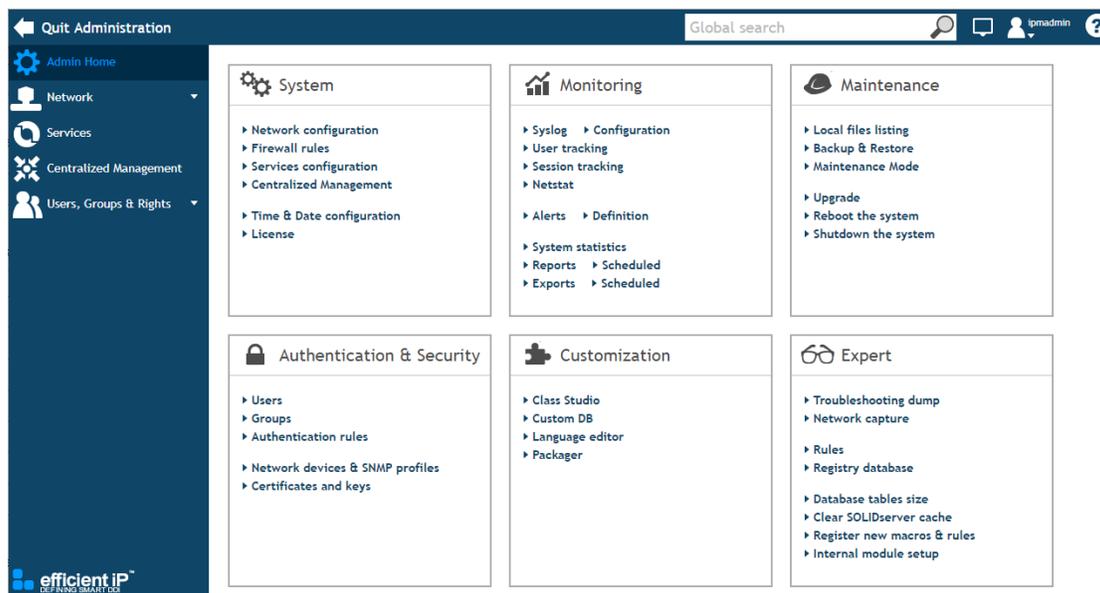


Figure 237. Admin Home, the homepage of the module Administration

Note that from the module **Dashboards**, you can gather gadgets and charts on *System dashboard* to monitor the module data or set up custom shortcuts and search engines. For more details, refer to the part [Dashboards](#).

This part contains the following chapters:

- [Centralized Management](#) describes how to manage remote SOLIDserver appliances, monitor them, configure their network and services and set up High Availability between two appliances.
- [Monitoring](#) describes how to monitor SOLIDserver from the reports, alerts, logs, statistics and users' sessions and performed operations to SNMP tools (profiles, monitoring details through the MIBs...).
- [Maintenance](#) describes how to maintain or troubleshoot SOLIDserver, manage local files, save/restore backup files or shut down and rebooting the appliance.
- [Upgrading](#) describes how to upgrade SOLIDserver, some checks must be performed before upgrading.

Some pages and options of the module Administration are detailed in other parts:

- The link *Certificates and keys* opens the pages *All certificates*, detailed in the section [Configuring the HTTPS Certificate](#), and *All GSS-TSIG keys*, detailed in the chapter [Implementing Dynamic Update](#).
 - The part [Configuring SOLIDserver](#) details *Network configuration*, *Services configuration* and *Licenses*.
 - The part [Rights Management](#) details *Users*, *Groups* and *Authentication rules*.
 - The part [Customization](#) details the GUI customization options, including *Language editor*, and the pages *Class Studio*, *Custom DB* and *Packager*. This part also details Smart Folders and IPv6 labels.
-

Chapter 90. Centralized Management

In the module Administration, the page *Centralized Management* allows administrators to:

- **Monitor appliances**, locally and remotely. The page returns system, hardware, licence and maintenance information. For more details, refer to the section .
- **Manage remotely** other appliances. From the Management appliance you can manage the service and network configuration, monitor and upgrade remote appliances.
- **Set up high availability** between the local appliance and a remote one. In such a configuration the Master appliance contains all the data you manage and the Hot Standby replicates the Master database but can have a specific configuration of its services and network.

Whether you want to manage remote appliances or set up a high availability configuration, you must:

1. **Configure the local appliance.**
2. **Add a remote appliance to the page Centralized Management.**

Note that with remote appliances or a high availability configuration, the upgrade can be performed remotely. For more details, refer to the chapter [Upgrading](#).

Browsing the Centralized Management Database

From the page **Centralized Management** you can configure and manage appliances remotely managed as well as the ones configured in High Availability (HA).

To access the page Centralized Management

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.

The appliances listed on the page do not have a properties page, all the information available can be displayed via the columns.

Understanding the Default Columns

The columns on the page provide details regarding the local appliance and any remote appliance, whether it is configured in high availability or not.

The **Role** allows to differentiate all appliances, they can be *Standalone*, *Master* or *Hot Standby*.

Table 90.1. The default columns on the page Centralized Management

Column	Description
Name	The appliance hostname. That name is used in the list <i>SOLIDserver</i> to differentiate local and remote appliances.
Local	Indicates which appliance you are connected to. Every appliance listed is either marked <i>Yes</i> (the appliance you are operating on) or <i>No</i> (the remote appliances).

Centralized Management

Column	Description
Manufacturer	The appliance manufacturer. It can therefore tell you if you are working on or remotely managing virtual appliances.
Product	The appliance model. It indicates the type of appliance you are working with and its size if it is not a virtual appliance (e.g. <i>SOLIDserver-250</i>).
Serial #	The appliance serial number. If two appliances have the same name, you can differentiate them using their serial number.
Version	The version of SOLIDserver: <i>current</i> is the latest.
IP address	The appliance IP address. The IP address of each appliance is used to access all the appliances that you want to add to the page <i>Centralized Management</i> .
Master address	The IP address of the Master appliance in the HA configuration. Master and Standalone appliances display <i>None</i> in this column.
Role	<p>Master: an appliance running in association with a Hot Standby appliance. It has a HA UID.</p> <p>Master (hot standby init): an appliance that used to be the Master in the HA configuration and is currently becoming the Hot Standby.</p> <p>Hot Standby: an appliance replicating the content of the Master appliance database it is associated with. It has the same HA UID as its master.</p> <p>Hot Standby (init): a Hot Standby appliance is being enrolled again with the same Master in case of replication failure.</p> <p>Standalone: an appliance configured and running on its own, with no HA configuration. This is the default role of any appliance.</p> <p>Standalone (hot standby init): an appliance becoming the Hot Standby of a Master appliance. It is not accessible for a few minutes, until the replication of the entire database is complete. During this time, the Hot Standby database is erased and replaced with the replication of the Master appliance database.</p> <p>Master (recovered): a Hot Standby appliance set as a Master is marked as such during the role switch, it is immediately operational.</p>
End of license	The date and time when a temporary license reaches its expiration. For appliances with an <i>official</i> license, this column is irrelevant.
End of maintenance	The date and time when the maintenance period reaches its expiration.
License time left	The number of days, hours and minutes until the <i>End of license</i> . For appliances with an <i>official</i> license, this column is irrelevant.
Maintenance time left	The number of days, hours and minutes until the <i>End of maintenance</i> .
System state	<p>The overall state of the appliance. It monitors and returns the information of the columns: <i>CPU load (5 min)</i>, <i>Disk I/O load (%)</i>, <i>Fans status</i>, <i>HDD space</i>, <i>LAGG status</i>, <i>Memory usage (%)</i>, <i>Power Supply Units status</i> and <i>RAID status</i>. These columns' statuses are described separately in the next table.</p> <p>OK: the appliance overall state is normal. None of the monitored columns is alerting.</p> <p>Unknown: one of the monitored columns returns no information.</p> <p>Critical: one of the monitored columns is in critical state.</p> <p>Warning: one of the monitored columns requires attention.</p>
HA UID	The key that identifies two appliances configured in High Availability.
Last write period	The last time the Hot Standby replicated the Master database.
Time drift	The difference, in seconds, between the Master NTP and the Hot Standby NTP. That drift should not exceed a minute (60 in the column) as this could have consequences on the DHCP failover replication.
Replication offset	The difference, in kilobytes, between the Master database and the Hot Standby database. As the replication is almost in real time, the difference should be minimal. A great value in this column could indicate a network disruption. If the Replication offset is <i>Unknown</i> , the remote SOLIDserver is in <i>Timeout</i> .

Column	Description
Status	OK: the appliance is up and running.
	Not configured: the local appliance has not been configured yet.
	Upgrading...: the Hot Standby appliance is being upgraded from the Master appliance.
	Switching to Hot Standby: an appliance is switching to Hot Standby role.
	Invalid credentials: the appliance has been switched to Standalone, was restored or the password of the SSH <i>admin</i> account has changed.
	Managed (remote): an appliance is being managed remotely, i.e. listed on the page Centralized Management of another appliance.
	Timeout: the appliance is not responding.
	Split-brain: two appliances are in Restricted mode due to a split-brain. For more details, refer to the section Troubleshooting a Split-brain .
Replication stopped: the connection is lost between the appliances. For more details, refer to the section Configuring Specific Behaviors if the Replication Takes a Long Time or Stopped .	

Understanding the Other Columns Available

In addition to the default columns, you can display extra columns dedicated, among other things, to system and hardware information. For more details regarding column display, refer to the section [Customizing the List Layout](#).

Table 90.2. The other columns available on the page Centralized Management

Column	Description
CPU load (5 min)	The load of all the appliances' CPU cores, on an average of 5 minutes. You can monitor and use the values in this column to return the specific statuses in the column <i>System State</i> , as detailed in the section Configuring Specific Thresholds to Monitor the Column System State .
Device OS version	The appliance architecture, either <i>amd64</i> or <i>i386</i> . <i>i386</i> is only displayed for remotely managed appliances in version prior to 6.0.x.
Disk I/O load (%)	The load of the appliance disk I/O, in percent. You can monitor and use the values in this column to return the specific statuses in the column <i>System State</i> , as detailed in the section Configuring Specific Thresholds to Monitor the Column System State .
Fans status	OK: the appliance fan(s) is/are up and running.
	N/A: no information on the appliance fan(s) is available.
	Disabled: the appliance fan(s) is/are set to disabled.
	Critical: at least one appliance fan returns an error (not running, on failure...).
Firmware date	The software image release date.
HDD space	The space used on the partition <i>/data1</i> of the appliance, in percent. You can monitor and use the values in this column to return the specific statuses in the column <i>System State</i> , as detailed in the section Configuring Specific Thresholds to Monitor the Column System State .
LAGG status	OK: the appliance LAGG interface(s) is/are up and running.
	N/A: no information can be retrieved from the appliance LAGG interface(s).
	Disabled: no LAGG interface is configured as active on the appliance.
	Critical: at least one appliance LAGG interface is down.
Last write time	The exact time of the last database replication.
Memory usage (%)	The memory usage of the appliance, in percent. You can monitor and use the values in this column to return the specific statuses in the column <i>System State</i> , as detailed in the section Configuring Specific Thresholds to Monitor the Column System State .

Column	Description
Multi-status	The multi-status of the appliance.
Power Supply Units status	OK: the appliance power supply unit(s) is/are up and running.
	N/A: no information can be retrieved from the power supply unit(s). For instance, on a virtual appliance or if the appliance has only one PSU.
	Disabled: no PSU redundancy is available on the appliance.
	Critical: one of the appliances' power supply units is unplugged or defective.
RAID status	The status of the RAID disk(s) of the physical appliances SDS-550, 1100, 2200, 3300 or Blast.
	OK: the appliance RAID disk(s) is/are up and running.
	N/A: no information can be retrieved from the appliance RAID disk(s).
	Disabled: no RAID disk is enabled on the appliance.
	Critical: the appliance RAID disk(s) is/are enabled but is degraded or offline.
Remote time	The time of the remote appliance(s) managed through the local appliance.
Time	The appliance date and time.

Configuring SOLIDserver to Remotely Manage Other Appliances

Whether you want to manage remote appliances or use a remote appliance to set up high availability, you must first of all configure the Management or Master SOLIDserver locally.

Configuring an appliance locally means assigning it an IP address. It sets the grounds for remote management and differentiates appliances:

- For remote management, it differentiates the Management appliance from the remote appliance.
- For high availability configurations, it differentiates the Master appliance from the Hot Standby appliance.

To configure locally the Management or Master appliance

Only users of the group *admin* can perform this operation.

1. Connect to the **future Master** or **Management** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page opens.

In the column **Local**, your appliance is marked  **Yes**. It does not have an IP address, which is why its **Status** is *Not configured*.

4. In the menu, select  **Tools** > **Configure local SOLIDserver**. The wizard **Configure local SOLIDserver** opens.
5. In the drop-down list **SOLIDserver IP address**, select the IP address of the appliance you are currently configuring.
6. Click on  to complete the operation. The report opens and closes. On the page *Centralized Management*, the local appliance details are now complete. Its **Role** is *Standalone* and its **Status**  **OK**.

From then on you simply need to add other appliances and remotely manage them from the list Centralized Management and the pages Network configuration and Services configuration.

Adding Remote Appliances

Once you configured your local appliance, you can add other appliances. You can add a large number of appliances to overview all SOLIDserver appliances used on your network.

The local appliance becomes a management platform where you remotely manage and/or monitor other SOLIDserver via the drop-down list SOLIDserver available on the pages *Network configuration*, *Services configuration*, *Syslog* and *System statistics* of the module Administration.

Prerequisites

- You must have at least two SOLIDserver appliances.
- You must configure the management platform locally, for more details refer to the section [Configuring SOLIDserver to Remotely Manage Other Appliances](#).
- The appliances remotely managed must be in software version 5.0.0.P0 or higher.
- The remote management can only be configured from and with appliances using an IPv4 address.
- On all appliances, the NTP should be configured to make sure they are all set at the same time and date. For more details, refer to the section [Configuring NTP Servers](#).

Adding a Remote Appliance

Once you configured locally the future Master or Management appliance, you can add remote appliances to the page *Centralized Management*. Keep in mind that:

- You can add as many remote appliances as you want to the page.
- You cannot display several High Availability configuration pairs on the page.
- Appliances remotely managed are still accessible locally.

To add a remote appliance

1. Connect to the **future Master** or **Management** appliance GUI.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. In the menu, click on **+ Add**. The wizard **Add/modify remote SOLIDserver** opens.
5. If you or your administrator created classes at the SOLIDserver (appliance) level, in the list **SOLIDserver (appliance) class**, select a class or *None*. Click on **NEXT**. The next page of the wizard appears.

Note that if no class is set and enabled, the class dedicated page of the wizard is automatically skipped. Applying a class on an object can impact the configuration fields available and/or required, for more information contact your administrator.

6. In the field **SOLIDserver IP address**, fill in the IP address of the appliance you want to add to the list.

7. In the field "**Admin**" **account password**, you can type in the password of your SSH account. By default, the default password of the account, *admin*, is automatically specified in the field.
8. If you want to monitor the remote appliance, tick the box **SNMP parameters**. The configuration fields open.

Table 90.3. SNMP parameters used to monitor the remote appliance

Field	Description
SNMP port	The port used to retrieve the remote appliance statistics. By default, the port 161 is used. If you changed the UDP port of your SNMP server, you must use the same port. For more details, refer to the section Managing the SNMP Service .
SNMP profile	The SNMP profile used to retrieve the statistics. By default, <i>standard v2c</i> is selected. The list contains the default profiles (<i>standard v1</i> , <i>standard v2c</i> and <i>standard v3</i>) and the ones you may have created. Each profile has its own level of security and enables the definition of a global security policy. For more details, refer to the section Managing SNMP Profiles .
SNMP retries	The number of connection attempts when the server is in timeout. You can set it between 0 and 5. By default, it is set to 2 attempts.
SNMP timeout	The number of seconds between each connection attempt. You can set it between 1s and 5s or set it to 10s. By default, it is set to 5.

9. Click on to complete the operation. The new appliance is listed. Its **Role** is *Standalone* and its **Status** is  *Remote (managed)*.

Once you added a remote appliance you monitor it, manage it or configure it in high availability. For more details, refer to the sections [Managing the Services and Network Configuration of Another Appliance](#), [Monitoring the Appliances Managed from the Centralized Management](#) and [Configuring Two Appliances in High Availability](#).

Managing the Services and Network Configuration of Another Appliance

Once you added an appliance to the page *Centralized Management*, you can remotely manage its services and network configuration from a dedicated drop-down list.

Managing the Network Configuration of Another Appliance

From the page *Network configuration* of a Management or Master appliance, you can manage the configuration of any appliance listed on the page *Centralized Management*.

To remotely manage the network configuration of another appliance

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, select the appliance of your choice. The page refreshes.

All appliances are listed as follows: *<hostname> (<IP address>)*. If an appliance is not reachable, it is listed as follows: *<hostname> (<IP address>) - Timeout*.

4. Edit the network configuration according to your needs. For more details, refer to the chapter [Configuring the Network](#).

Managing the Services Configuration of a Remote Appliance

From the page *Services configuration* of a Management or Master appliance, you can manage the configuration of any appliance listed on the page *Centralized Management*.

Keep in mind that **you can edit any service except the source email address of the alert**. The address *noreply@efficientip.com* that sends you the alert notifications has to be edited locally.

To remotely manage the services configuration of another appliance

1. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, select the appliance of your choice. The page refreshes.

All appliances are listed as follows: *<hostname> (<IP address>)*. If an appliance is not reachable, it is listed as follows: *<hostname> (<IP address>) - Timeout*.

4. Edit the service configuration according to your needs. For more details, refer to the chapter [Configuring the Services](#).

Monitoring the Appliances Managed from the Centralized Management

Once you added remote appliances to the page *Centralized Management* you can monitor their logs, their statistics and any time drift between the local appliance and the remote appliances.

Monitoring the Logs of Another Appliance

From the page *Syslog*, you can monitor the logs of all the appliances you manage remotely.

All the remote management and High Availability related logs respect the format: *HA <event>*.

To monitor the logs of a remote appliance

All the remote management and High Availability related logs respect the format: *HA <event>*.

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** or **Management** appliance GUI.
2. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
4. Under the menu, in the drop-down list **SOLIDserver**, select the remote appliance of your choice. The page refreshes.

All appliances are listed as follows: *<hostname> (<IP address>)*. If an appliance is not reachable, it is listed as follows: *<hostname> (<IP address>) - Timeout*.

5. In the drop-down list **Service**, select *ipmserver*.

6. In the search engine of the column **Log**, type in *HA ** and hit **Enter**. Only the remote management and high availability logs are returned.

Monitoring the Statistics of Another Appliance

From the page *System statistics* you can display the charts of the appliances you manage remotely.

Note that, if your remote appliances are in version 6.0.1 or higher, you can even generate the statistics reports of a remote appliances. For more details, refer to the section [Generating a Report](#).

To access the Statistics of a remote SOLIDserver

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **System statistics**. The page **System statistics** opens.
3. Under the menu, in the drop-down list **SOLIDserver**, select the remote appliance of your choice. The page refreshes to display the statistics of the appliance you selected.

You can only display the statistics of a remote appliance in version 6.0.1 or higher, appliances in lower versions are listed as follows: *<hostname> (<IP address>) - Not supported*.

For more details regarding each chart, refer to the section [Monitoring the Appliance Statistics](#).

Monitoring the Time and Date Drift Between Appliances

On the page *Centralized Management*, the column *Time drift* allows to monitor and compare the time and date of your local and remote appliances.

With high availability configurations, monitoring any time drift between the Master and Hot Standby is paramount. Both appliances should be set at the same time. To ensure there is no shortage of data on the Hot Standby appliance in case it needs to become a Master:

1. **We strongly recommend that you configure the time and date of both the Master and Hot Standby appliances via NTP servers.** For more details, refer to the section [Configuring NTP Servers](#).
2. You can configure the alert *Member clock drift* that monitors the time drift between the Master and the Hot Standby.

If you do not want to enroll any remote appliance as Hot Standby, the synchronization is optional. If your local and remote appliances are set at the exact same time, it is useless to configure the alert *Member clock drift*; you might even disable it. For more details, refer to the section [Enabling or Disabling Alerts](#).

Configuring the Alert that Monitors the Time Drift

The alert *Member clock drift* monitors the time difference between the Management or Master appliance and the remote appliance(s) it manages. The alert is triggered if the time of any remote appliance, whether a Hot Standby or not, has drifted too much from the management appliance time.

You can edit the alert definition settings to suit your needs. By default:

- The alert is based on a filter of the column *Time drift*. The filter is set with a 20 seconds threshold as follows: $> 20 \parallel < -20$. The alert is raised:
 - If the absolute time difference between the remote appliance and the Management or Master appliance is higher than 20 seconds.
 - If the absolute time difference between the remote appliance and the Management or Master appliance is lower than 20 seconds.

To edit the threshold, refer to the procedure [To edit the threshold of the alert Member clock drift](#).

- The time difference check is performed every 5 minutes. To edit the frequency, refer to the procedure [To edit the alert Member clock drift default frequency and recipients](#).
- The alert sends an email alert to the users of the group *admin*. If you did not configure the users' email address properly, they are not notified. To edit the recipients, refer to the procedure [To edit the alert Member clock drift default frequency and recipients](#) below.

To edit the threshold of the alert Member clock drift

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** or **Management** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
4. Filter the list if need be.
5. In the column **Alert filter**, click on *Edit alert filters*. You are redirected to the page **Centralized Management**, where the alert was created. Above the menu, a blue banner indicates you are in *Alert edition mode*.
6. In the search engine of the column **Time drift**, the default filter $> 20 \parallel < -20$ is visible.
7. Edit the filter according to your needs.
8. On the right-end side of the menu, click on . The wizard **Quit editing the alert filters** opens.
9. Tick the box **Save changes before quitting** to save your new filter.
10. Click on to complete the operation. The report opens and closes. You are redirected back to the page **Alerts Definition**.

To edit the alert Member clock drift default frequency and recipients

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** or **Management** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
4. At the end of the line of the *Member clock drift*, click on . The properties page opens.
5. In the panel **Main properties**, click on . The wizard **Edit an alert definition** opens.
6. Tick the box **Expert mode**.
7. To edit the alert check frequency:

- a. Tick the box **Edit scheduling**. The schedule drop-down lists appear.
 - b. By default, only the drop-down list *Minutes* is set to *Every 5 minutes*.
 - c. Select the values that suit your needs.
8. To edit the group of users receiving the alert or add recipients:
- a. Make sure the box **Send mail** is checked.
 - b. In the drop-down list **Mailing lists**, select the group of users of your choice.
 - c. In the field **Additional Mail**, type in the email address of a user that does not belong to the group selected. Click on **ADD**. The email address is moved to the **Additional Mail List**. Repeat these actions for as many users as you need.
9. Click on **OK** to complete the operation. The report opens and closes. The properties page is visible again.

Configuring Specific Thresholds to Monitor the Column System State

The column **System state** returns the overall state of the appliance based on the information returned by the columns *CPU load (5 min)*, *Disk I/O load (%)*, *Fans status*, *HDD space*, *LAGG status*, *Memory usage (%)*, *Power Supply Units status* and *RAID status*.

You can edit some registry database entries to set up thresholds and ensure that the system state is *Warning* or *Critical* based on the value of the columns *CPU load (5 min)*, *Disk I/O load (%)*, *HDD space* and *Memory usage (%)*.

Table 90.4. The configurable registry database entries to return Critical and Warning System states

Registry database entry	Description
module.system.member_snmp_cpu_crit	Returns <i>Critical</i> if the value of the column <i>CPU load (5 min)</i> matches or exceeds the value you set for the entry. By default, it is set to 150.
module.system.member_snmp_cpu_warn	Returns <i>Warning</i> if the value of the column <i>CPU load (5 min)</i> matches or exceeds the value you set for the entry. By default, it is set to 100.
module.system.member_snmp_hdd_crit	Returns <i>Critical</i> if the value of the column <i>HDD space</i> matches or exceeds the value you set for the entry. By default, it is set to 90 %.
module.system.member_snmp_hdd_warn	Returns <i>Warning</i> if the value of the column <i>HDD space</i> matches or exceeds the value you set for the entry. By default, it is set to 80 %.
module.system.member_snmp_ios_crit	Returns <i>Critical</i> if the value of the column <i>Disk I/O load (%)</i> matches or exceeds the value you set for the entry. By default, it is set to 50 %.
module.system.member_snmp_ios_warn	Returns <i>Warning</i> if the value of the column <i>Disk I/O load (%)</i> matches or exceeds the value you set for the entry. By default, it is set to 25 %.
module.system.member_snmp_mem_crit	Returns <i>Critical</i> if the value of the column <i>Disk I/O load (%)</i> matches or exceeds the value you set for the entry. By default, it is set to 100 %.
module.system.member_snmp_mem_warn	Returns <i>Warning</i> if the value of the column <i>Disk I/O load (%)</i> matches or exceeds the value you set for the entry. By default, it is set to 98 %.

The thresholds that you set only apply to the column System State on the appliance where you configure the registry database. For that reason, if you manage remote appliances or a high availability configuration, you should edit the entries on the Management or Master appliance.

To set up a threshold to return a Warning or Critical state in the column System State

Only users of the group *admin* can perform this operation.

1. If you manage remote appliances or a high availability configuration, connect to the **Master** or **Management** appliance GUI.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
4. In the column **Name**, type in *module.system.member_snmp*. The matching entries are returned, for more details refer to the [table](#).
5. In the column **Value**, click on the line of your choice. The wizard opens.
6. In the field **Value**, specify the value that suits your needs.
7. Click on **OK** to complete the operation. The new Value is visible on the page.

Configuring Two Appliances in High Availability

The High Availability (HA) is a system network design that ensures that your network continues to respond even if one or more of its components fail. This architecture provides integrated disaster recovery management operations for transparent and efficient service continuity. It also prevents you from losing any data if anything were to happen to your managing platform.

With SOLIDserver, High Availability implies that you connect together two appliances where one local appliance is a Master and the other, a remote one, is a Hot Standby, i.e. a read-only backup server replicating the content of the Master appliance database.



Figure 90.1. High Availability representation

The Master and Hot Standby appliances work together to make sure that **when the automatic switch is enabled, if the Master crashes or encounters any problem, the Hot Standby can replace it immediately and vice versa**. Therefore, the Hot Standby must replicate the Master database as often as possible.

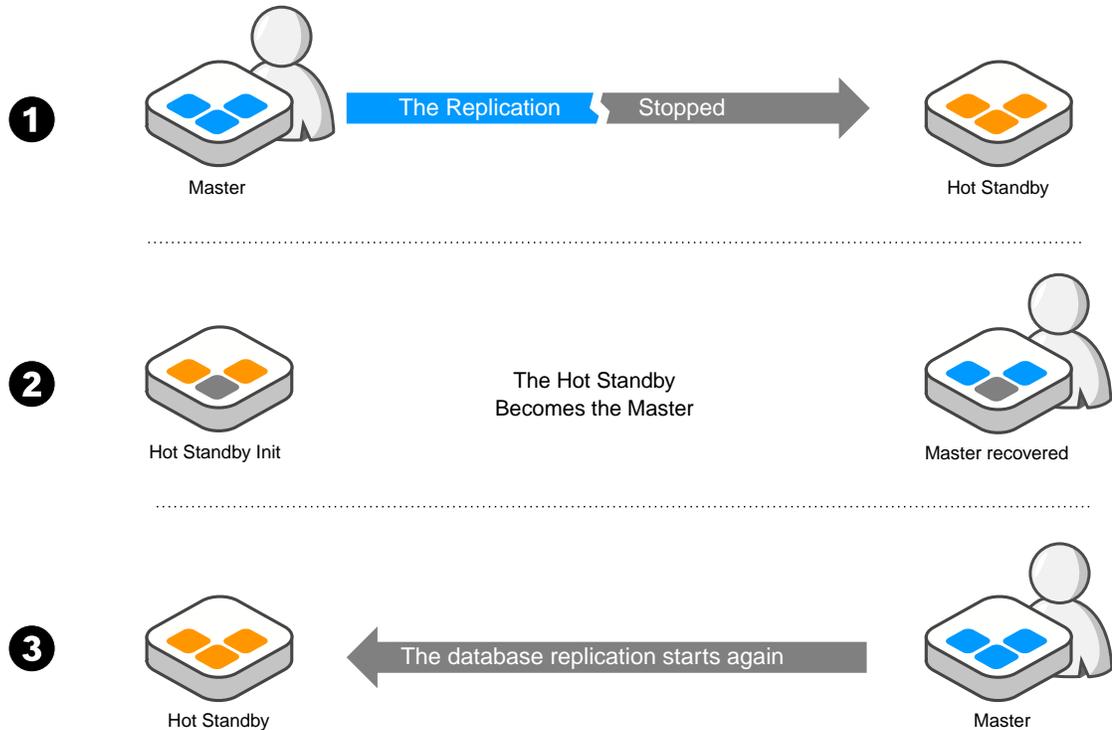


Figure 90.2. Switch mechanism if the replication stops

Prerequisites

- You must have two SOLIDserver appliances.
- The HA configuration can only be configured from and with appliances using an IPv4 address.
- The Master appliance should be configured locally as detailed in the section [Configuring SOLIDserver to Remotely Manage Other Appliances](#).
- The future Hot Standby must be added to the page Centralized Management of the Master as detailed in the section [Adding Remote Appliances](#).

Limitations

- The database High Availability is configurable **only for two appliances**.
- **We strongly advise against displaying several HA configurations on the page Centralized Management.** If you add an appliance to this list, it means that you want to manage it. Therefore, if you decide to add to your managing appliance two appliances configured in High Availability, it means that you intend to manage them from the managing appliance. On the page Centralized Management of the appliances in HA, the appliance Status changes from *OK* to *Invalid credentials* because the local *admin* management password overwrites the management password locally set on the Master appliance of this other HA configuration.
- **The HA does not support the configuration of a NAT between the two appliances.** Both appliances send their local IP address when they communicate, therefore the converted IP address cannot be used in the HA communication. Configuring a NAT might even break the HA configuration.

Setting a High Availability Configuration

Once the Master appliance is configured locally and the future Hot Standby is added to the page *Centralized Management* of the Master, you can configure high availability between the appliances, i.e. enroll the Hot Standby.

This configuration has to be done from the future Master appliance and can be done on layer 2 or 3 or the network. For more details, refer to the section [Frequently Asked Questions](#).

Keep in mind that for the configuration to be viable and effective **the two appliances must:**

- Match the [prerequisites](#).
- Be set at the **same time**. For more details, refer to the section [Configuring NTP Servers](#).
- Have the **same version** of SOLIDserver.
- Have the **same performance** rate, to ensure a smooth transition. In the event of a switch, the former Hot Standby has retrieved all the database information and can actually provide the same performance and efficiency as the original Master.
- Have the **same architecture** (32 bits or 64 bits).

To configure High Availability between two appliances

Only users of the group *admin* can perform this operation.

1. Connect to the **future Master** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. Tick the appliance you want to set up as Hot Standby.
5. In the menu, select  **Edit** > **Enroll SOLIDserver as Hot Standby**. The wizard **Enroll SOLIDserver as Hot Standby** opens.
6. Click on to complete the operation. The Report opens and works for a while until the Hot Standby appliance database is erased and replaced by the Master appliance database. The appliance set as Hot Standby is unavailable for a while. Each appliance *Role* is modified according to the configuration, they now share the same *HA UID*.

Once the High Availability is configured:

1. The information on the page *Centralized Management* of the Master is now also available on the Hot Standby appliance.
2. On the page *Centralized Management* of the Hot Standby appliance, the only operation available is switching the appliances role. For more details, refer to the section [Switching the High Availability Configuration](#).
3. The Hot Standby appliance is now in read-only mode. Every modification made on the Master appliance is copied in the Hot Standby database almost in real-time.

You can no longer edit the remote appliance database locally but you can still edit the services and network configuration of both appliances separately. For more details, refer to the sections [Managing the Network Configuration of a Remote Appliance](#) and [Managing the Services Configuration of a Remote Appliance](#).

4. The Hot Standby appliance replicates the content of the Master appliance database to provide an efficient backup if it has to replace the current Master appliance.
 - From the page *Centralized Management* of the Master appliance, you should monitor the columns **Time drift** and **Replication offset**, to make sure that the Hot Standby appliance properly replicates the database. If at some point the replication stops, you can enroll again the Hot Standby appliance following the procedure [To configure High Availability between two appliances](#).
 - You can, for instance, configure and enable the automatic switch so that, if the Hot Standby has not replicated the Master database in the last 60 seconds, it should check the Master status three times in a row, every 4 seconds. If there is no response (timeout, etc.), the Hot Standby switches to Master. For more details, refer to the section [Controlling the Automatic Switch Mechanisms if the Network is Unreliable](#).

Note that **the automatic switch is not enabled by default**. You can manually enable it. For more details, refer to section [Configuring High Availability Advanced Options](#).

Editing Remote Appliances

You can edit the local and remote appliance(s) IP address, *admin* account password and/or SNMP monitoring parameters.

To edit a remote appliance

1. Connect to the **future Master** or **Management** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. In the column **Name**, right-click on the appliance of your choice. The contextual menu opens.
5. Click on  **Edit**. The wizard **Add/modify remote SOLIDserver** opens.
6. Edit the fields **SOLIDserver IP address**, **"Admin" account password** and **SNMP parameters** according to your needs. For more details, refer to the section [Adding a Remote Appliance](#).

Note that for security purposes, the field *"Admin" account password* is emptied out in the edition wizard.

7. Click on to complete the operation. The new appliance is listed. Its **Role** is *Standalone* and its **Status** is  *Remote (managed)*.

Managing a High Availability Configuration

Once you have successfully configured appliances in high availability, you can:

- **Monitor** the Hot Standby from the Master appliance. For more details, refer to the section [Monitoring the Appliances Managed from the Centralized Management](#).
- **Manage** the configurations on the pages *Services configuration* and *Network configuration* of the Hot Standby from the Master appliance. For more details, refer to the section [Managing the Services and Network Configuration of Another Appliance](#).
- **Update the database of the Hot Standby**, as detailed in the section [Updating the Database of the Hot Standby](#).

- **Configure advanced options** to make sure your high availability appliances behave as expected if your network is unreliable or may be disrupted or if the replication is slow or stopped, as detailed in the section [Configuring High Availability Advanced Options](#).
- **Replace SOLIDserver appliances** and keep a viable high availability configuration. You can replace either appliance, with or without backup, as detailed in the section [Replacing An Appliance in High Availability](#).
- **Troubleshoot the configuration** as detailed in the section [High Availability Configuration Troubleshooting Solutions](#).
- **Answer some frequently asked questions**, all detailed in the section [High Availability Configuration Troubleshooting Solutions](#).

Updating the Database of the Hot Standby

At any time, you can execute the option *Update HA files database* to update the database of the Hot Standby.

This option is useful if any new file on the Master appliance has not been replicated yet on the Hot Standby or if you intend to manually switch the appliances roles in the configuration, as detailed in the section [Switching the High Availability Configuration](#)

Note that if you enabled the automatic switch, the option is automatically launched before the switch.

To update the database of appliances in High Availability

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. In the menu, select  **Tools > Update HA files database**. The report opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again and the columns **Last write periods**, **Time drift** and **Replication offset** are updated.

Switching the High Availability Configuration

Switching the high availability configuration means switching the *Role* of each appliance in the configuration, the Master becomes the Hot Standby and vice versa.

Via the option *Manually switch local SOLIDserver to master* you can switch a configuration. Note that:

- The option **must be executed from the Hot Standby** as you cannot make a Master change its role to Hot Standby.
- The **switch is no longer automatic by default**. If you want to automatically detect potential problems on the Master (timeout, crash...) and switch the appliances roles, you must enable this behavior. For more details, refer to the section [Configuring High Availability Advanced Options](#).
- Even with the automatic switch enabled, you can manually switch the configuration at any time.

To switch the appliances role in a high availability configuration

Only users of the group *admin* can perform this operation.

1. Connect to the **Hot Standby** appliance GUI. The message **This SOLIDserver is a Hot Standby: Database is in READ-ONLY mode** is present on every page.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. Update the Hot Standby database:
 - a. In the menu, select  **Tools > Update HA files database**. The report opens.
 - b. Click on **OK** to complete the operation. The report opens and closes. The page is visible again and the columns **Last write periods**, **Time drift** and **Replication offset** are updated.
5. Switch the configuration:
 - a. In the menu, select  **Tools > Manually switch local SOLIDserver to master**. The report opens.
 - b. Click on **CLOSE** to commit the changes. The page **Centralized Management** is visible.

The **Role** of the former Hot Standby appliance is **Master (recovered)**. The former Master appliance is marked **Master (Hot Standby init)**. For more details, refer to the table [The default columns on the page Centralized Management](#).

The Hot Standby appliance is unavailable until it has replicated the Master database

Configuring High Availability Advanced Options

Some registry database keys allow you to customize the High Availability switch and automatic re-enrollment mechanisms.

Checking if the Automatic Switch is Enabled

Once you configured the high availability, you can set up an automatic switch of the appliances role. First of all, you need to check if the behavior is enabled.

To check if the automatic switch is enabled

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
4. In the search engine of the column **Name**, type in *module.system.max_hot_standby_time_skew*. If the value is *0* or *-1*, the automatic switch is disabled. To enable it or set custom switch behaviors, refer to the section [Controlling the Automatic Switch Mechanisms if the Network is Unreliable](#).

Controlling the Automatic Switch Mechanisms if the Network is Unreliable

If your network is unreliable, you can monitor both appliances thanks to:

- The column **Offset replication**, it allows to monitor the Hot Standby database reliability.
- Some keys in the registry database that allow to control the switch mechanism, avoid flapping and ensure that the switch occurs if and only if there is a problem on the Master appliance.

The High Availability dedicated keys each serve a specific purpose and can impact or be impacted by one another. They allow to:

Control the automatic switch

The key `module.system.max_hot_standby_time_skew` allows to control and potentially prevent the automatic switch itself, based on the appliances *Last write period*.

By default, its value is set to 3600 seconds, if the last write period is greater than an hour, the two appliances cannot switch automatically. **This key is disabled on fresh installations and the switch is no longer automatic by default**, its value is -1. You can edit the key to enable it again. Note that after an upgrade from a previous version, it can still be enabled.

The key can be set with values between -1 and 2^{31} . Setting the key to 0 or -1 prevents the automatic switch altogether, regardless of the values set for the keys described below.

Keep in mind that if you stop the automatic switch, if your appliances stop communicating and the Hot Standby stopped replicating the Master database, you must configure the HA again. For more details, refer to the section [Disabling the HA Configuration Without Losing the Database](#).

Set the maximum period of time a switch should take

The key `module.system.init_hot_standby_timeout` allows to control how long a switch should take, whether you are enrolling an appliance or switching the roles.

By default, it is set to 1000 seconds, if the Hot Standby initialization stage has not evolved after 16 minutes, the enrollment or switch stops and the appliance keeps its current role as Standalone or Master.

Setting a high value for this key is useful if the Master database is very large or if your network is not reliable.

Set a replication lag period before switching

The key `module.system.hot_standby_max_replication_lag` allows to set a lack of database replication period before automatically switching the appliances role.

By default, it is set to 60 seconds, if the Hot Standby has not replicated the database in the last 60 seconds, it tries to contact the Master appliance n times (n is defined by the key `module.system.hot_standby_switch_retry`). If the Master is not responding, i.e. it does not send its role and status, the Hot Standby switches to Master.

Set the number of retries before switching

The key `module.system.hot_standby_switch_retry` allows to control the number of retries before automatically switching the appliances role.

By default, it is set to 3 attempts, if the Hot Standby appliance cannot connect to the Master and check its role and status, it tries to get an answer 3 times in a row. If after 3 attempts there is still no answer, it takes over the Master role.

Keep in mind that the retries check frequency is defined by the key `module.system.hot_standby_switch_sleep` detailed below.

Set a connection timeout between the appliances

The key `module.system.hot_standby_connect_timeout` allows to set a maximum response time for the Master appliance.

By default, it is set to 4 seconds. If after 4 seconds, the connection is not established, the Hot Standby considers the Master appliance to be down. It then waits for n seconds (the value of the key `module.system.hot_standby_switch_sleep`) until it tries establishing a connection again.

Keep in mind that the number of connection retries is defined by the key `module.system.hot_standby_switch_retry`, and that the retry frequency depends on the key `module.system.hot_standby_switch_sleep`.

Set the frequency of the retries

The key `module.system.hot_standby_switch_sleep` allows to set the retries check frequency.

By default, it is set to 4 seconds, if the Hot Standby does not get an answer from the Master, it tries every 4 seconds n times (depending on the number of retries you set for the key `module.system.hot_standby_switch_retry`).

Prevent unexpected switches during minor issues on the network

Among other configurations, the registry database entries can be used to prevent unexpected switches if the network experiences minor issues. To do so you should increase the switch lag period (the value of the key `module.system.hot_standby_max_replication_lag`) and/or the number of retries (the value of the key `module.system.hot_standby_switch_retry`). Keep in mind that a large number of retries might overload the network.

To disable the automatic switch altogether, refer to the section detailing how to [control the automatic switch](#).

To configure the high availability switch behavior from the registry database

Only users of the group `admin` can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in the name the key of your choice. For more details, refer to the [high availability keys](#).
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in the value of your choice. For more details, refer to the [high availability keys](#).
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Anticipating a Network Disruption

At some point you might plan on disrupting the network or shutting it down (for potential repairs, equipment changes, etc.), in this case, **we strongly recommend that you disable the HA configuration**.

Indeed, in the event of a network disruption, voluntary or not, the appliances configured in HA go in Timeout or switch their roles but might not successfully go through with it. In this case, when the appliances start again you might have two Masters. The original Master keeps its role and the Hot Standby switches to *Master (recovered)*, with no Hot Standby configured and potentially a case of *split-brain*. If it is the case, there is no accurate way of knowing which appliance should become the Hot Standby. For more details, refer to the section [Troubleshooting a Split-brain](#).

To prevent any loss of data, we suggest that you follow the procedure in the section [Disabling the High Availability Configuration](#). Once the network is back on, you must configure the HA again.

Configuring Specific Behaviors if the Replication Takes a Long Time or Stopped

A set of keys in the registry database allow to **automatically re-enroll the Hot Standby**, that is to say replicate again the entire database of the Master. You can customize the triggers or disable the functionality. The automatic re-enrollment applies if:

The database replication stopped

If the replication stopped because the Hot Standby is no longer connected, the Master automatically re-enrolls it. On the page *Centralized Management*, the Hot Standby status is *Replication stopped*.

The key `module.system.auto_replication_repair` is enabled by default. If you want to disable the automatic re-enrollment, you must set the value of the key to `0`.

The database replication is taking a long time

With a large *Replication offset* a switch could make you lose the latest changes. You can monitor the offset and set a threshold to automatically re-enroll your Hot Standby.

- The key `module.system.warning_replication_lag` allows to set the maximum replication offset before generating a warning message in the logs. By default, it is set to `10240` kB.
- The key `module.system.max_replication_lag` allows to set the maximum replication offset before re-enrolling the Hot Standby. By default, it is set to `307200` kB. Note that if the automatic re-enrollment is disabled, the offset is ignored.

To make sure that the Master does not try re-enrolling the Hot Standby indefinitely, it automatically tries the re-enrollment a maximum of 3 times over 24 hours. If it does not succeed, a message is displayed in the GUI and you need to troubleshoot the HA yourself. For more details, refer to the section [Troubleshooting the Messages](#).

- The key `module.system.auto_replication_repair.threshold` allows to set how long the maximum replication offset can be exceeded before re-enrolling the Hot Standby. By default, it is set to `60` times 10 seconds, i.e. if the `module.system.max_replication_lag` is exceeded during 10 minutes or more, the re-enrollment is triggered.

Note that to be notified of any replication delay before the Hot Standby is automatically re-enrolled, the value of the key `module.system.warning_replication_lag` should be lower than the value of the key `module.system.max_replication_lag`.

To disable the automatic re-enrollment of the Hot Standby

Only users of the group `admin` can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.

3. In the search engine of the column **Name**, type in *module.system.auto_replication_repair*. Only this key is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in *0* to disable it. The default value is *1*.

Note that once the key is disabled, the key *module.system.max_replication_lag* is ignored.

6. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

To set the maximum replication offset before generating a log message

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the search engine of the column **Name**, type in *module.system.warning_replication_lag*. Only this key is listed.
4. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in the value of your choice, in kB. The default value is *10240*.
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

To set the maximum replication offset before re-enrollment

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. Make sure the key *module.system.auto_replication_repair* is enabled, set to *1*. If it is disabled, the automatic re-enrollment is not triggered.
3. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
4. In the search engine of the column **Name**, type in *module.system.max_replication_lag*. Only this key is listed.
5. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
6. In the field **Value**, type in the value of your choice, in kB. The default value is *307200*.
7. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

To set how long the maximum replication offset can be exceeded before re-enrollment

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. Make sure the key *module.system.auto_replication_repair* is enabled, set to *1*. If it is disabled, the automatic re-enrollment is not triggered.

3. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
4. In the search engine of the column **Name**, type in `module.system.auto_replication_repair.threshold`. Only this key is listed.
5. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.
6. In the field **Value**, type in the value of your choice. The default value is `60` time 10 seconds, i.e. 10 minutes.
7. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

High Availability Configuration Troubleshooting Solutions

You need to troubleshoot the HA configuration if there is a replication problem or if you have an improper configuration. Keep in mind that **the configuration is viable and effective only if the two appliances:**

- **Are set at the same time.** For more details, refer to the section [Configuring NTP Servers](#).
- **Have the same version of SOLIDserver.**
- **Have the same performance rate**, to ensure a smooth transition. In the event of a switch, the former Hot Standby has retrieved all the database information and can actually provide the same performance and efficiency as the original Master.
- **Have the same architecture** (32 bits or 64 bits).

If you do not meet all the requirements listed above, warning messages are displayed on the Hot Standby appliance GUI.

Troubleshooting the Messages

To help you monitor the High Availability configuration a set of red messages can be displayed on all the pages of the appliances to inform or warn you.

This SOLIDserver is a Hot Standby: Database is in READ-ONLY mode.

- *Appliance:* Hot Standby.
- *Problem:* There is no configuration problem, this message indicates that you cannot edit the Hot Standby appliance database locally, it automatically replicates the Master database.
- *Solution:* N/A. This message is informative.

This SOLIDserver is initializing the Master appliance database replication.

- *Appliance:* Hot Standby, when it is still in *Hot Standby (init)*.
- *Problem:* There is no problem, the replication is starting. The Hot Standby database should soon be in read-only.
- *Solution:* Wait until the replication is complete. This message is informative.

This SOLIDserver cannot automatically switch to Master: its database is out-of-date.

- *Appliance:* Master recovered.

- *Problem:* The replication offset is higher than the limit set in the key `module.system.max_hot_standby_time_skew`. Therefore during the appliances switch, the Hot Standby became Master recovered and the former Master never became Hot Standby.
- *Solution:* Disable the High Availability configuration following the split-brain [Manual Resolution](#).

High Availability Configuration: impossible to reach the remote SOLIDserver.

- *Appliance:* Master.
- *Problem:* The Hot Standby is unreachable, so it does not replicate the Master database.
- *Solution:* Delete the Hot Standby from the page Centralized Management, add it to the list again and enroll it as detailed in the section [Configuring Two Appliances in High Availability](#).

The two SOLIDserver appliances you configured in High Availability have DIFFERENT ARCHITECTURES. They must have the same architecture for the database replication to work.

- *Appliance:* Master.
- *Problem:* The Master and Hot Standby appliances do not have the same architecture version, therefore the database replication is impossible.
- *Solution:* Upgrade the Hot Standby appliance from the Master appliance: the Hot Standby is automatically upgraded to the architecture version of the Master. For more details, refer to the section [Upgrading Appliances Managed Remotely](#) in the chapter [Upgrading](#).

The two SOLIDserver appliances you configured in High Availability have DIFFERENT VERSIONS. To avoid configuration problems you should upgrade the Hot Standby from the Master.

- *Appliance:* Master.
- *Problem:* The Master and Hot Standby appliances do not have the same version of SOLIDserver, therefore the database replication may encounter some problems.
- *Solution:* Upgrade the Hot Standby appliance from the Master appliance: the Hot Standby is automatically upgraded to the version of SOLIDserver of the Master. For more details, refer to the section [Upgrading Appliances Managed Remotely](#) in the chapter [Upgrading](#).

The information used to identify the local SOLIDserver has changed.

- *Appliance:* Master and Standalone.
- *Problem:* The manufacturer, product and/or serial number of the Master appliance¹ have been edited or you restored another appliance's backup. These three pieces of information are needed to identify the appliance.
- *Solution:* Delete the unwanted remote appliances from the page Centralized Management and configure the high availability again following the procedures in the sections [Configuring Two Appliances in High Availability](#).

A case of SPLIT-BRAIN was detected: both SOLIDserver appliances configured in HA are set as Master. As changes have been made through both servers, you are in Restricted mode. To go back to a proper HA configuration you need to choose which SOLIDserver should remain the Master: connect as an administrator to the Master and on the page Centralized Management select Tools > Manually switch this SOLIDserver to master.

¹This information is displayed in the gadget *System information* of the Master appliance Home page.

- *Appliance:* Master and Hot Standby.
- *Problem:* During the appliances switch, the Hot Standby became Master and the former Master stayed Master as well. The automated detection could not resolve the situation. For more details, refer to the section [Troubleshooting a Split-brain](#).
- *Solution:* Disable the High Availability configuration following the split-brain [Manual Resolution](#).

You are in Restricted mode. From the page Centralized Management you can go back to Normal mode: either configure the local SOLIDserver or delete all the remote appliances.

- *Appliance:* Master and Hot Standby.
- *Problem:* Either your appliances are in split-brain, one of them is in timeout or the replication stopped.
- *Solution:* Follow the troubleshooting solution of the other message displayed in the GUI. Each problem has a specific solution.

The Hot Standby stopped database replication stopped.

- *Appliance:* Master.
- *Problem:* The Hot Standby can no longer replicate the Master database. If they switch, you lose the latest changes performed on the Master appliance.
- *Solution:* Find the origin of the problem, you might need to disable the High Availability and configure it again following the procedures in the sections [Configuring Two Appliances in High Availability](#).

Troubleshooting a Split-brain

The Split-brain is a very specific case that might occur when two appliances are configured in High Availability. Once configured, both appliances share the same HA UID but in the event of a network disruption they might end up sharing the same role instead of keeping the two roles on which the configuration relies: one Master and one Hot Standby.

This HA configuration requires that the Master appliance database can be edited and that its Hot Standby replicates it, in real time, to be ready for a potential switch of the configuration. With two Masters, whether they are both Master or one is Master and the other is Master (recovered), there is no backup and both appliances can potentially overwrite each other's changes.

To help you in the prevention of the Split-brain, SOLIDserver follows a set of checks, when the appliances communicate once again:

1. SOLIDserver starts up in Restricted mode and runs in normal mode if and only if no HA conflicts are detected.
2. SOLIDserver checks if both appliances share the same version. If not, a message is displayed on every page of the appliance with the latest version.
3. SOLIDserver checks if both appliances share the same role.

If it turns out that both appliances are Master, there is a set of resolutions that SOLIDserver tries and executes on its own to avoid staying in Restricted mode.

Automated Detection

When a Master appliance detects that the other appliance is also a Master, SOLIDserver performs three checks to try and avoid a case of split-brain:

1. **If no appliance has been edited since the last synchronization:** the last appliance that switched to Master remains Master and enrolls the other appliance as Hot Standby.
2. **If one appliance has been edited since the last synchronization:** the last appliance that was modified becomes Master and enrolls the other appliance as Hot Standby.
3. **If both appliances have been edited since the last synchronization:** SOLIDserver puts them in Restricted mode with the status *Split-brain*, as specified in the red message displayed on every page of both appliances. To configure the HA again, you have to execute a [Manual Resolution](#) as detailed in the section below.

Manual Resolution

The manual resolution is only needed when the appliances in HA are in a case of split-brain that puts them in Restricted mode, as specified in the red message displayed on every page of both appliances. This mode implies two behaviors:

1. The synchronization between the appliances stopped, as if you had two Standalone appliances with the same HA UID.
2. You can still edit the database of both appliances but no changes are actually pushed on the physical server(s).

To go back to a viable configuration, you have two possibilities:

- Disable the High Availability and configure it again as described in the section [Disabling the HA Configuration Without Losing the Database](#).
- Force the configuration and choose which appliance becomes the Master as described in the section [Switching the High Availability Configuration](#).

Disabling the HA Configuration Without Losing the Database

In some cases, like in the event of a split-brain, you might want to disable the HA configuration but keep both appliances database. In this case, you need to switch one appliance to Standalone.

Before switching an appliance to Standalone, keep in mind that:

1. You should only switch an appliance to Standalone if:
 - Your appliances in High Availability are in Split-brain. In which case, you end up with two Master appliances and need to reconfigure the HA. For more details, refer to the section [Troubleshooting a Split-brain](#).
 - Your appliances in High Availability disconnected somehow and the automatic switch was disabled using the registry key detailed in the section [Controlling the Automatic Switch Mechanisms if the Network is Unreliable](#). In which case, the replication stopped and you need to set it up again.

In any other case, disabling the configuration should be done using the standard procedure detailed in the section [Disabling the High Availability Configuration](#).

2. It is impossible to switch a Master appliance to Standalone if it has a Hot Standby.
3. The Hot Standby appliance must always be the first one to be switched to Standalone.
4. Switching an appliance to Standalone erases the database entirely whether the appliance is a Hot Standby or a Master. So if you switch both appliances to Standalone, you erase your entire database.
5. Switching an appliance to Standalone automatically saves a backup file for the appliance.

6. Switching an appliance to Standalone has to be done locally. You must connect to the appliance via its IP address.

To disable HA configuration by putting the Hot Standby appliance in standalone

Only users of the group *admin* can perform this operation.

1. Connect to your **Hot Standby** appliance GUI.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. In the menu, select **Tools > Reset local SOLIDserver to standalone**. The wizard **Reset local SOLIDserver to standalone** opens.
5. Click on **OK** to complete the operation. The report opens and works for a while, a backup of the database is saved before the appliance database is erased.

The appliance **Role** is now *Standalone*. During the switch, the appliance might be unavailable.

Once the former Hot Standby appliance has rebooted and is reachable:

- Connect to its GUI via its IP address and log in using the credentials of a user belonging to the group *admin*.
- Configure the *Internal module setup* once again. For more details, refer to the section [Defining the Internal Module Setup](#).
- Make sure the appliance is a Standalone, either in the General information gadget of the appliance home page or on the page Centralized Management directly.
- You can display or download its backup file from the page *Backup & Restore*. For more details, refer to the section [Managing Backups and Restoring Configurations](#).

To reconfigure the high availability, you must:

1. Connect to the Master appliance using its IP address.
2. Go to the page Centralized Management and make sure the former Hot Standby appliance has:
 - The Role *Standalone*.
 - The status *OK*.
3. If the Standalone appliance status is OK, you must:
 - Enroll it as the Hot Standby. For more details, refer to the procedure [To configure High Availability between two appliances](#).
4. If the Standalone appliance status is not OK, you must:
 - Delete it from the page Centralized Management of the Master appliance. For more details, refer to the procedure [To disable the high availability configuration](#).
 - Add it to list the Centralized Management again. For more details, refer to the procedure [To add a remote appliance](#).
 - Enroll it as the Hot Standby. For more details, refer to the procedure [To configure High Availability between two appliances](#).

Frequently Asked Questions

1. Can I be warned of configuration problems?

Yes, you can set an alert on the page Centralized Management to be informed of any change via email or SNMP trap. For more details regarding the alerts configuration, refer to the chapter [Managing Alerts](#).

For instance, filter the appliances through the *Status* column to detect a split-brain or any status different from OK. If your configuration ends up in split-brain, refer to the section [Troubleshooting a Split-brain](#).

2. Can I set up the HA on any network layer?

The appliances can be configured on layer 2 or 3 of the network.

Layer 2 configuration: If the appliances are configured on layer 2, they belong to the same LAN. Therefore you can set up a VIP interface that would allow you to access the current Master appliance of the configuration through the IP address you set (the original master if it is acting as a master, or the Hot Standby if the configuration was switched). For more details, refer to the section [Configuring a VIP](#) in the chapter [Configuring the Network](#).

Layer 3 configuration: If the appliances are configured on layer 3, they do not belong to the same LAN. The HA is still configurable and running perfectly through the routers that connect them but it is impossible to set a VIP to access the Master appliance.

3. Can I customize the HA mechanism to suit my network?

Yes, if your network is unreliable or experiences frequent disruptions, your administrator can follow the procedure of the section [Configuring High Availability Advanced Options](#). It details the registry database key that you can edit to modify the HA number of retries or configure automatic switch parameters.

4. What can I do from the Hot Standby appliance?

You cannot edit the database as it is replicating the content of the Master database. However, from the Hot Standby appliance you can:

- Switch the appliances roles and convert the Hot Standby in Master. For more details, refer to the section [Switching the High Availability Configuration](#).
- Switch the Hot Standby to *Standalone*. This operation must be performed carefully as it erases the database entirely and disables the High Availability configuration. This operation is only recommended in the event of a split-brain. For more details, refer to the section [Disabling the HA Configuration Without Losing the Database](#).
- Set up the Network configuration page according to your needs. This page is independent from the database and can therefore be configured differently on the Master and Hot Standby appliances. For more details, refer to the chapters [Centralized Management](#) and [Configuring the Network](#).
- Set up the System configuration page according to your needs. This page is independent from the database and can therefore be configured differently on the Master and Hot Standby appliances. For more details, refer to the chapters [Centralized Management](#) and [Configuring the Network](#).
- Save a backup of the appliance. For more details, refer to the section [Managing Backups and Restoring Configurations](#).

However, you cannot restore the backup of an appliance in High Availability. You need to disable the High Availability, restore the backup and then configure the High Availability again.

Replacing Appliances Managed Remotely

At some point, you might need to replace an appliance that you manage from the page Centralized Management.

To replace appliances you must take into account if they are simply remotely managed or configured in high availability.

Replacing an Appliance Managed Remotely

To successfully replace an appliance managed remotely you need to:

1. **Remove the appliance from the page Centralized Management of the managing appliance**

For more details, refer to the section [Deleting Appliances Managed Remotely](#).

2. **Add the new appliance to the page Centralized Management of the managing appliance**

- First, you need to configure locally the new appliance. For more details, refer to the section [Configuring SOLIDserver to Remotely Manage Other Appliances](#).
- Second, you need to add the new appliance to the page Centralized Management of the managing appliance. For more details, refer to the section [Adding Remote Appliances](#).

Replacing An Appliance in High Availability

With appliances configured in high availability, **to prevent any loss of data, you must always replace the Hot Standby appliance**. There are two scenarios possible:

1. You can replace an appliance that has a backup, following the recommendations of the section [Replacing a Hot Standby Appliance With Backup](#).
2. You can replace an appliance that has no backup available, following the recommendations of the section [Replacing a Hot Standby Appliance Without Backup](#).

Replacing a Hot Standby Appliance With Backup

If you generated a backup of the appliance you need to replace, you must follow the steps below.

1. **If the appliance that needs to be replaced is the Master, switch its role to Hot Standby.** For more details, refer to the section [Switching the High Availability Configuration](#).
2. **Disable the High Availability configuration** from the Master appliance: delete the Hot Standby from the page *Centralized Management*. The Hot Standby switches to *Standalone*. For more details, refer to the section [Disabling the High Availability Configuration](#).
3. **Restore the backup on the Standalone appliance.** For more details, refer to the procedure [To restore a backup file](#).
4. **Switch the current Master to Standalone** only once when the restoration of the backup on the other appliance is complete, this ensures your database availability. You now have two Standalone appliances: one with backup that becomes the Master, and one totally empty that becomes the Hot Standby.

5. Configure the High Availability again:

- First, connect to the appliance where you restored the backup, the future Master appliance, and add the future Hot Standby to the page Centralized Management. For more details, refer to the section [Adding a Remote Appliance](#).
- Second, enroll the Hot Standby. For more details, refer to the procedure [To configure High Availability between two appliances](#) in the section Configuring HA Management.

6. If you restored your backup directly on your Hot Standby, manually switch the appliances' role.

If you followed this procedure from step 2, your current Master used to be the Hot Standby so you may need to switch their role back again. For more details, refer to the section [Switching the High Availability Configuration](#).

Replacing a Hot Standby Appliance Without Backup

The replacement of an appliance in HA with no backup must follow the steps below:

1. **Put the appliance that needs to be replaced in Hot Standby role**, if it is currently the Master. For more details, refer to the section [Switching the High Availability Configuration](#).
2. **Disable the High Availability configuration**. For more details, refer to the section [Disabling the High Availability Configuration](#).
3. **Set the network and services configuration of the future Hot Standby appliance according to your needs**. For more details, refer to the sections [Configuring the Network](#) and [Configuring the Services](#).

We strongly recommend that you use an NTP server to set both appliances at the time.

4. **Add the new appliance to the page Centralized Management of the Master appliance and enroll it:**
 - First, you need to add the new appliance to the page Centralized Management of the Master appliance. For more details, refer to the section [Adding a Remote Appliance](#)
 - Second, you need to enroll the new appliance as Hot Standby. For more details, refer to the procedure [To configure High Availability between two appliances](#).
5. **Manually switch the configuration if the new appliance is supposed to be the Master in the configuration**. For more details, refer to the section [Switching the High Availability Configuration](#).

Deleting Remote Appliances

At any point you can delete a remote appliance from the page *Centralized Management*.

Before going further keep in mind that:

- It is impossible to delete the local appliance.
- For appliances remotely managed, **deleting an appliance allows to stop managing it remotely**. For more details, refer to the section [Deleting Appliances Managed Remotely](#).
- For appliances in high availability, **deleting the Hot Standby disables the high availability**. For more details, refer to the section [Disabling the High Availability Configuration](#)

Deleting Appliances Managed Remotely

Deleting a remote appliance removes it from the list. You no longer manage the appliance.

If the appliance is configured in high availability, refer to the section [Disabling the High Availability Configuration](#).

To delete an appliance from the page **Centralized Management**

1. Connect to the **Master** or **Management** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. Tick the appliance(s) of your choice.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on to complete the operation. The report opens and closes. The appliance is not listed anymore.

Disabling the High Availability Configuration

From the page *Centralized Management* of the Master appliance you can delete the Hot Standby. Before going further, keep in mind that:

- **Deleting the Hot Standby disables the configuration** because its *Role* is revoked, the communication with the Master breaks and therefore the replication stops.
- **The database of the Hot Standby is erased.**

If you want to disable the High Availability and keep the Hot Standby database intact, refer to the section [Disabling the HA Configuration Without Losing the Database](#).

- **The HA UID of the Master appliance changes** for two reasons:
 1. To prevent any other appliance from managing it as it would delete its database.
 2. The HA UID can be used again during the next HA configuration with this appliance as a Master.

To disable the high availability configuration

Only users of the group *admin* can perform this operation.

1. Connect to the **Master** appliance GUI.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
4. Tick the Hot Standby appliance.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on to complete the operation. The report opens and works for a while - the Hot Standby database is saved before it is erased - and closes. The Hot Standby is no longer listed on the page. The former Master appliance keeps a *Master* role.

The former Hot Standby appliance is not accessible while it is reset.

Once the former Hot Standby has rebooted and is reachable:

- Connect to the Hot Standby via its IP address, log in using the credentials of a user belonging to the group *admin*.
- The appliance is no longer in read-only mode but its database is empty.
- Configure the *Internal module setup* once again. For more details, refer to the section [Defining the Internal Module Setup](#).
- The page *Centralized Management* displays the following information:
 1. The appliance itself (*Local*) is the only one listed, the former Master appliance is no longer part of that list.
 2. The appliance role is now *Standalone*.
 3. The appliance needs to be configured locally again (Tools > Configure local SOLIDserver).

Chapter 91. Monitoring

There are many ways of monitoring SOLIDserver, for administrators and standard users, from the resources to the logs and SNMP profiles and more. This chapter gathers the following monitoring possibilities:

- [Managing Reports](#).
- [Managing Alerts](#).
- [Managing the Logs](#).
- [Monitoring the Appliance Statistics](#).
- [Tracking Sessions](#).
- [Tracking Users](#).
- [Managing SNMP Profiles](#).
- [Monitoring Using SNMP](#).
- [Displaying Netstat Data](#).
- [Sizing the Database Tables](#).

Note that from the page **Centralized Management**, you can monitor the status, hardware or license and maintenance information of an appliance, local or remote. For more details, refer to the section [Centralized Management](#).

Managing Reports

From the modules DHCP, DNS, NetChange and Administration you can generate a number of reports to *HTML* or *PDF* format. These reports can be generated at a given time or scheduled. All existing reports are detailed in the section [Browsing the Reports Database](#).

To export your data, refer to the chapter [Exporting Data](#).

Browsing the Reports Database

To list the reports

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Reports**. The **Reports** page opens.

To list the scheduled reports

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Reports*, click on **Scheduled**. The page **Scheduled reports** opens.

DHCP Reports

You can generate advanced reports for DHCP servers and scopes.

DHCP Server Reports

All the DHCP server dedicated reports are available in v4 and v6 on the page *All servers*, on the server properties page and on all the listing pages displaying the content of a specific server. To generate these reports, refer to the section [Generating a Report](#).

Clients Most Used OS

- *Prerequisite:* Selecting at least one DHCP server managing IPv4.
- *Description:* Contains a table, for each physical server selected, detailing the lease clients most used OS for all the scopes of the server. It contains the OS name, the percentage of leases that use it and the total number of leases that it corresponds to.

Server Data Exchanges

- *Prerequisite:* Selecting at least one DHCP server.
- *Description:* Contains charts detailing the physical servers data exchanges (*Discover*, *Offer*, *Request* and *Acknowledge*) over the *last 24 hours*, *last 7 days*, *last 30 days* and *last 12 months*.

Server Options Comparison

- *Prerequisite:* Selecting at least two DHCP servers.
- *Description:* Compares one by one all the DHCP options configured on the selected servers. For more details regarding DHCP options, refer to the chapter [DHCP Options](#).

Server Usage Chart

- *Prerequisite:* Selecting at least one DHCP server.
- *Description:* Contains charts detailing the physical servers IP addressing management, number of allocated leases, number of used statics, number of free IP addresses: over the *last 24 hours*, *last 7 days*, *last 30 days* and *last 12 months*.

Most Used Networks

- *Prerequisite:* Selecting at least one DHCP server managing IPv4.
- *Description:* Contains a table for each selected physical server detailing the most used scopes and what they contain.

Server Usage Evolution Charts

- *Prerequisite:* Selecting at least one server.
- *Description:* Contains lease and request dedicated charts providing an overview of a server usage evolution. The chart results are based on server usage a daily, monthly, semestrial and yearly basis.

DHCP Scope Reports

All the DHCP scope dedicated reports are available in v4 on the page *All scopes*, on the scope properties page and on all the listing pages displaying the content of a specific scope. To generate these reports, refer to the section [Generating a Report](#).

Clients Most Used OS

- *Prerequisite:* Selecting at least one DHCP server managing IPv4.

- *Description:* Contains a table, for each physical server scope selected, detailing the lease clients most used OS. It contains the OS name, the percentage of leases that use it and the total number of leases that it corresponds to.

Scopes Options Comparison

- *Prerequisite:* Selecting at least two scopes.
- *Description:* Compares one by one all the DHCP options configured on the selected scopes. For more details regarding DHCP options, refer to the chapter [DHCP Options](#).

Scopes Summary

- *Prerequisite:* Selecting at least one scope.
- *Description:* Provides detailed tables of the DHCP options activity and origin of the selected scope(s). For instance, it indicates if the option was set at scope level or inherited from the managing server.

DNS Reports

You can generate advanced reports for DNS servers, views and zones.

DNS Server Reports

All the DNS server dedicated reports are available on the page *All servers*, on the server properties page and on all the listing pages displaying the content of a specific server. To generate these reports, refer to the section [Generating a Report](#).

A records without an IPv4 address or alias

- *Prerequisite:* Selecting at least one server configured with DNS to IPAM replication.
- *Description:* Contains the list of A records that are not associated to an IPv4 address or alias, by server.

AAAA records without an IPv6 address or alias

- *Prerequisite:* Selecting at least one server configured with DNS to IPAM replication.
- *Description:* Contains the list of AAAA records that are not associated to an IPv6 address or alias, by server.

CNAME records without an alias

- *Prerequisite:* Selecting at least one server configured with DNS to IPAM replication.
- *Description:* Contains the list of CNAME records that are not associated to an IPv4 or IPv6 alias, by server.

PTR records without an IPv4 address

- *Prerequisite:* Selecting at least one server configured with DNS to IPAM replication.
- *Description:* Contains the list of PTR records that are not associated to an IPv4 address, by server.

PTR records without an IPv6 address

- *Prerequisite:* Selecting at least one server configured with DNS to IPAM replication.

- *Description:* Contains the list of PTR records that are not associated to an IPv6 address, by server.

Route 53 Incompatibilities

- *Prerequisite:* Selecting at least one Amazon Route 53 server.
- *Description:* Contains a list of all the selected server options and configurations incompatible with Amazon Route 53. No piece of information listed in the tables can ever be replicated on the Amazon Route 53 server you are managing via SOLIDserver.

Zones NS and IP addresses

- *Prerequisite:* Selecting at least one Amazon Route 53 server.
- *Description:* Contains a list of all the NS records and their corresponding IP address of all the zones of the server. It also indicates if they were pushed to the smart architecture, that is to say replicated on all the servers managed by the smart.

Servers Configuration

- *Prerequisite:* Selecting at least one server.
- *Description:* Contains all the server configuration details divided into 4 tables: *Settings* (all the options), *ACLs* (all the access control lists), *Keys* (all the DNS keys configured) and *Groups* (all the group of users that have access to the server).

Hybrid DNS Engine Incompatibilities

- *Prerequisite:* Selecting at least one physical server.
- *Description:* Contains the list of all the options and configuration that make the server incompatible with Hybrid. For more details, refer to the section [Generating the Hybrid Incompatibilities Report](#).

Server Peak Hour

- *Prerequisite:* Selecting at least one server; except an Amazon Route 53 server. Cloud servers cannot be monitored by this report.
- *Description:* Contains a table displaying the selected server peak hour over the last 24 hours. It indicates the number of queries per second processed by the server at its busiest hour in the last 24 hours.

Servers Configuration Comparison

- *Prerequisite:* Selecting at least two servers.
- *Description:* Contains tables that allow to compare the selected servers configurations: *DNS server parameters*, *DNS server ACLs* and *DNS server keys*.

Query Rate per Server

- *Prerequisite:* Selecting at least one server; except an Amazon Route 53 server. Cloud servers cannot be monitored by this report.
- *Description:* Contains four charts representing the response to queries curve of the selected server: over the *last 24 hours*, *last 7 days*, *last 30 days* and *last 12 months*.

Server Reply to Queries Charts

- *Prerequisite:* Selecting at least one server; except an Amazon Route 53 server. Cloud servers cannot be monitored by this report.
- *Description:* Contains four charts representing the four types of response to queries (*success, failure, inexistent RR set and Nxdomain*) for the selected server: over the *last 24 hours, last 7 days, last 30 days* and *last 12 months*.

Server Usage Charts

- *Prerequisite:* Selecting at least one server.
- *Description:* Contains usage evolution charts for the selected server: queries over the past week, last 6 months, last month and past year at the time of the generation of the report.

DNS View Reports

All the DNS view dedicated reports are available on the page *All views*, on the view properties page and on all the listing pages displaying the content of a specific view. To generate these reports, refer to the section [Generating a Report](#).

View Statistics

- *Prerequisite:* Selecting at least one view.
- *Description:* Contains two tables for every selected view. The first table displays the number of zones and indicates how many are Forward zones. The second table details all the zones of the view with: their name, their type and the amount of known¹ records they contain.

DNS Zone Reports

All the DNS zone dedicated reports are available on the page *All zones*, on the zone properties page and on all the listing pages displaying the content of a specific zone. To generate these reports, refer to the section [Generating a Report](#).

Route 53 Incompatibilities

- *Prerequisite:* Selecting at least one Amazon Route 53 server.
- *Description:* Contains a list of all the selected zone options and configurations incompatible with Amazon Route 53. No piece of information listed in the tables can ever be replicated on the Amazon Route 53 server you are managing via SOLIDserver.

Zones NS and IP addresses

- *Prerequisite:* Selecting at least one Amazon Route 53 zone.
- *Description:* Contains a list of all the NS records and their corresponding IP address for the selected zone. It also indicates if they were pushed to the smart architecture, that is to say replicated on all the servers managed by the smart.

Zones Missing RRs

- *Prerequisite:* Selecting at least one zone.
- *Description:* Contains a list of all the misconfigured records within the selected zones divided into 5 tables: *PTR records without A or AAAA, A or AAAA records without PTR, CNAME records without A or AAAA, NS records without A or AAAA* and *MX records without A or AAAA*.

¹By *known* records, we mean all the record types that have ever been created or managed on the appliance.

Zone Statistics

- *Prerequisite:* Selecting at least one zone.
- *Description:* Contains tables listing all the selected zones: name, type and amount of known² records they contain. The report contains as many tables as there are managing servers among the selected zones.

Zones Configuration Comparison

- *Prerequisite:* Selecting at least two zones.
- *Description:* Contains tables detailing the *allow-transfer*, *allow-update*, *forward*, *masters* and *notify* parameters configuration for the selected zone(s). Each parameter value is listed with the zone name it is configured for and the server it belongs to.

NetChange Reports

You can generate advanced reports for NetChange network devices.

NetChange Network Device Reports

All NetChange dedicated reports are available on the page *All network devices*, on the network device properties page and on all the listing pages displaying the content of a specific network device. To generate these reports, refer to the section [Generating a Report](#).

Network Devices Properties

- *Prerequisite:* Selecting at least one device.
- *Description:* Contains basic information regarding the selected device(s): the *Device name*, *Device type*, *Ports usage (%)* and *Ports used*.

NetChange/IPAM/DHCP data comparison

- *Prerequisite:* No need to select any device. The report automatically takes into account all the devices.
- *Description:* Contains data discrepancies between NetChange and other modules. All the inconsistencies found are listed in 10 tables focusing on IPv4 and IPv6 addressing. These tables compare NetChange and IPAM data at device and at discovered items level (IP/MAC addresses association) as well as NetChange and DHCP data (MAC addresses and lease association).

Network devices summary

- *Prerequisite:* No need to select any device. The report automatically takes into account all the devices.
- *Description:* Contains information regarding all the network devices you manage through SOLIDserver divided into four sections: *Summary* that contains all the network devices dedicated pie charts available by default on NetChange dashboard, *Network devices model by vendor* that contains charts displaying your devices vendors and models, *Top 50 most used network devices* that contains a table listing the most used devices with the percent of port usage and the total number of used ports and finally *Top 50 most unused network devices* that contains a table listing the least used devices with the percent of port usage and the total number of used ports.

²By *known* records, we mean all the record types that have ever been created or managed on the appliance.

Administration Reports

From the module Administration, you can generate advanced reports for the appliance statistics and users.

Appliance Statistics Reports

All the appliance statistics dedicated reports are available on the page *System statistics*. To generate these reports, refer to the section [Generating a Report](#).

Statistics chart

- *Prerequisite*: No selection required. However, only users of the group *admin* can access the page.
- *Description*: Contains all the charts available on the page *System statistics*. Their content depends on the time of the generation.

Network traffic

- *Prerequisite*: No selection required. However, only users of the group *admin* can access the page.
- *Description*: Contains charts representing all the ingoing and outgoing traffic on your network over the last 24 hours, last 7 days, last 30 days and last 12 months.

User Reports

A report allows to export all the permissions of a user from the page *Users*. To generate this report, refer to the section [Generating a Report](#).

Users rights in each group

- *Prerequisite*: Selection at least one user.
- *Description*: Contains a table displaying for the selected user(s) their *User name*, *Group name*, *Resources* and *User name*. In other words, the group(s) of users they belong to, the objects they have access to and the actions they can perform on these objects.

The report is empty for disabled users.

Generating a Report

From a set of listing and properties pages you can generate and download reports in PDF or HTML format.

Keep in mind that if the report you are generating contains charts, choosing the *HTML* format requires you to be connected to the appliance via its domain name when you open the report, otherwise the charts are empty.

To generate a report from a listing page

1. Go to the page of your choice. All existing reports are detailed in the section [Browsing the Reports Database](#).
2. If necessary, tick the object(s) you want to generate the report for.

Note that from the page **System statistics**, in the drop-down list **SOLIDserver**, you can select your local appliance or a remote one, only if the remote appliance is in version 6.0.1 or higher.

3. In the menu, select  **Report** > <report-of-your-choice> ³. The corresponding wizard opens.
4. In the list **Report format**, select an export format, either **HTML** or **PDF**. By default, *HTML* is selected.
5. Click on **NEXT**. The next page of the wizard opens.
6. In the drop-down list **Action**, select **Generate new data**. If you already have generated a report for the same object, the drop-down list allows to select and generate it again.
7. Click on **OK** to generate the report. The page *Report* opens and works for a while.
8. You can click on **DOWNLOAD** to save the report immediately.

When the report is generated, it is available on the page *Reports*. For more details, refer to the procedure [To list the reports](#).

9. Click on **CLOSE** to close the wizard.

To generate a report from a properties page

1. Go to the page of your choice. All existing reports are detailed in the section [Browsing the Reports Database](#).
2. At the end of the line of the object of your choice, click on . The properties page opens.
3. In the menu, select  **Report** > <report-of-your-choice> ⁴. The corresponding wizard opens.
4. In the list **Report format**, select an export format, either **HTML** or **PDF**. By default, *HTML* is selected.
5. Click on **NEXT**. The next page of the wizard opens.
6. In the drop-down list **Action**, select **Generate new data**. If you already have generated a report for the same object, the drop-down list allows to select and generate it again.
7. Click on **OK** to generate the report. The page *Report* opens and works for a while.
8. You can click on **DOWNLOAD** to save the report immediately.

When the report is generated, it is available on the page *Reports*. For more details, refer to the procedure [To list the reports](#).

9. Click on **CLOSE** to close the wizard.

Scheduling a Report

The generation of reports can easily be scheduled for all types of reports through the same wizard as for immediate generation.

To schedule the generation of a report

1. Go to the page of your choice. All existing reports are detailed in the section [Browsing the Reports Database](#).
2. If necessary, tick the object(s) you want to generate the report for.

³The reports dedicated to *Compare DNS data with IPAM data* and *Amazon Route 53* are all listed in the submenu of the same name.

⁴The reports dedicated to *Compare DNS data with IPAM data* and *Amazon Route 53* are all listed in the submenu of the same name.

3. In the menu, select **Report** > <report-of-your-choice>. The corresponding wizard opens.
4. In the list **Report format**, select an export format, either **HTML** or **PDF**. By default, *HTML* is selected.

Keep in mind that if the report you are generating contains charts, choosing the *HTML* format requires you to be connected to the appliance via its domain name when you open the report, otherwise the charts are empty.

5. Click on **NEXT**. The next page of the wizard opens.
6. In the drop-down list **Action**, select **Schedule the report**. The page refreshes and displays the scheduling fields.
7. Configure the export frequency or date and time of export using the table below.

Table 91.1. Scheduled report fields

Field	Description
Day(s) of the week	In this drop-down list, select a frequency (over the whole week or for a specific set of days) or a specific day of the week. By default, <i>Every day</i> is selected.
Date of the month	In this drop-down list, select a specific day of the month or a frequency (every day) for the refresh. By default, <i>Every day</i> is selected.
Month	In this drop-down list, select a specific month or <i>Every month</i> . By default, <i>Every month</i> is selected.
Hour	In this drop-down list, select a frequency (period of time), i.e. a set of hours or a specific hour of the day. By default, <i>20</i> is selected.
Minute	In this drop-down list, select a moment (o'clock, quarter past, half past or quarter to). By default, <i>00</i> is selected.
Name	In this field, name the scheduled export in this field. By default, it is named after the reports generation service. This field is required.
Mail to	In this drop-down list, select the group which users should receive the export notification email. This email cannot be sent if the users email address is not valid or if your SMTP relay is not configured. For more details, refer to the section Configuring the SMTP Relay . By default, the first of your groups, in the ASCII alphabetic order, is selected.
Rights as	In this drop-down list, select a user, his/her rights and limitations are applied to the report: only the items this user has access to are listed in the export.

8. Click on **OK** to complete the operation. The report opens and closes.

The scheduled report configuration is available on the page *Scheduled reports*, refer to the procedure [To list the scheduled reports](#).

When the report is generated, it is available on the page *Reports*. For more details, refer to the procedure [To list the reports](#).

For more details regarding scheduled reports refer to the section [Managing Scheduled Reports Configuration Files](#).

Downloading and Displaying Reports

From the page *Reports*, you can download PDF reports and open the HTML reports in a new tab of your browser.

All the generated reports are listed on this page whether they were generated at a specific time or scheduled.

To download a PDF report

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Reports**. The page **Reports** opens.
3. In the column **Format**, filter the list through to display only PDF reports.
4. In the column **Name**, click on the report of your choice to download it to your computer.

To open an HTML report

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Reports**. The page **Reports** opens.
3. In the column **Format**, filter the list through to display only PDF reports.
4. In the column **Name**, click on the report of your choice to download it to your computer.

Managing Scheduled Reports Configuration Files

The Scheduled reports page gathers the configuration details of all the scheduled reports. It allows you to disable and enable back a scheduling or delete it. Every time a report is generated, it is listed on the Reports page. For more details, refer to the section [Downloading and Displaying Reports](#).

To enable/disable a scheduled report

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Reports*, click on **Scheduled**. The page **Scheduled reports** opens.
3. All the scheduled reports are listed by name, report type and format.
4. Tick the scheduled report of your choice.
5. In the menu, select  **Edit** > **Enable** or **Disable**. The wizard opens.
6. Click on to complete the operation. The report opens and closes. The scheduled report configuration is now *OK* or *Disabled*.

To delete a scheduled report

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Reports*, click on **Scheduled**. The page **Scheduled reports** opens.
3. All the scheduled reports are listed by name, report type and format.
4. Tick the scheduled report you want to delete.
5. In the menu, click on  **Delete**. The **Delete** wizard opens.
6. Click on to complete the operation. The report opens and closes. The scheduled report configuration is no longer listed.

Managing Alerts

SOLIDserver offers a number of customization options that include the alert configuration from any page. You can be notified of the changes of your choice (new value, status, etc.) either via email or via an SNMP trap. Alerts provide an extra monitoring system.

Prerequisites

- To properly set an SNMP trap on an alert, make sure the SNMP and SMTP servers are properly configured. For more details, refer to the chapters [Managing the SNMP Service](#) and [Configuring the SMTP Relay](#).
- To properly set mail notifications on an alert, you need to specify a group of users or specific mail addresses. Make sure the email addresses of the group members and/or those you specified in the wizard are valid as incorrect email addresses cannot receive alerts. Also, make sure the groups you specify is configured with sufficient rights to assess the situation. For more details, refer to the chapter [Managing Users](#).

Browsing Alerts

The Administration module contains two pages dedicated to alerts. You cannot configure the columns display on these pages.

- The page **Alerts** displays the details of all the alerts that have been raised: their priority, when they were raised and released, their current state, etc.
- The page **Alerts Definition** contains all the alerts configured in SOLIDserver. It displays the configuration details of each alert, provides a link to edit the alert filters and allows to enable/disable each alert.

Browsing the Alerts Database

To display the page Alerts

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Alerts**. The page **Alerts** opens. Only the raised alerts are displayed by default.
3. To display all the alerts, whether they are dismissed or not, under the menu, tick the box **Display all alerts** .

To display the page Alerts Definition

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
3. To display the alerts related to a specific alert definition, in the column **Alert name**, click on the name of your choice. The page **Alerts** opens and displays only the alerts related to that specific definition.
4. To display all the alerts, whether they are dismissed or not, under the menu, tick the box **Display all alerts** .

To display an alert properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Alerts**. The page **Alerts** opens.
3. To display an alert properties page from the page **Alerts**, in the column **Alert**, click on the alert of your choice. The properties page opens.
4. To display an alert properties page from the page **Alerts Definition**:
 - a. In the breadcrumb, click on **Alerts Definition**. The page **Alerts Definition** opens.
 - b. At the end of the line of the alert name of your choice, click on . The properties page opens.

Understanding the Columns on the page Alerts Definition

By default, the page *Alerts Definition* provides an overview of all the alerts created and their configuration through six columns:

Table 91.2. The columns on the page Alerts Definition

Column	Description
Alert Name	The alert name.
Condition	The trigger condition of the alert, <i>0</i> being the status or value that triggers the alert if met (=), different (!=, <, >), etc.
Created on	The alert creation date and time.
Scheduling	The check frequency for the parameters that trigger the alert.
Alert filters	The filters set on the page before creating the alert. Each alert provides the link to edit these filters. For more details, refer to the section Editing the Alert Filters .
State	<i>Released</i> : once the alert is created it is <i>Released</i> on the page Alerts Definition. <i>Raised</i> : once the configured parameters are met, the alert is <i>Raised</i> and is listed on the page Alerts.
Status	The alert definition status, either <i>Enabled</i> or <i>Disabled</i> .

The alert properties page displays extra information that can be configured such as the severity, priority, recipients of the email, etc.

Understanding the Columns on the page Alerts

By default, the page *Alerts* displays all the raised alerts and their details.

The raised alerts are the ones that have not been dismissed or have been manually set back to monitored. To display all the formerly raised alerts, refer to the procedure [To display the page Alerts](#).

Table 91.3. The columns on the page Alerts

Column	Description
Severity	The alert severity that you set upon creation: <i>Minor</i> , <i>Major</i> , <i>Crash</i> or <i>Block</i> .
Module	The name of the module where the alert was created.
Sub Module	The name of the page within the module where the alert was created.
Alert Name	The name given to the alert upon creation.
Priority	The level of priority you set upon creation: <i>Low</i> , <i>Normal</i> , <i>High</i> , <i>Urgent</i> or <i>Immediate</i> .

Column	Description
Begin date	The time and date when the alert went from released to raised, i.e. the moment the parameters set were met.
End date	The time and date when the alert went from raised to released.
Starting since	The period of time since the alert has been raised.
Monitoring	<i>Monitored</i> : when you have not dismissed an alert or if you have reinstated it.
	<i>Dismissed</i> : when you have dismissed an alert.
State	<i>Released</i> : when the alert has reached its <i>End date</i> .
	<i>Raised</i> : when the parameters set for the alert are met, the alert is raised.

Adding Alerts

From any page within SOLIDserver you can create alerts from the menu  **Alerts, gadgets & Smart Folders**. Before adding alerts you can filter the list to customize the trigger and create the alerts that suit your needs. So if you decide to filter the page *All zones* via the column *Status* with *!=OK* and then add an alert, the alert would be triggered when any zone listed changes status to a status different from *OK* and send you an email and/or an SNMP trap depending on what you configure.

To add an alert

This procedure is an example, it sends an alert if any zone status changes to anything but *OK*.

1. Go to the page of your choice and filter the list according to your needs.
 - a. In the sidebar, go to  **DNS > Zones**. The page **All zones** opens.
 - b. In the column **Server**, click on the name of the server of your choice to display the zones it contains.
 - c. Double-click in the search engine of the column **Status**. The filter constructor appears.
 - d. In the drop-down list, select *!= (different from)*.
 - e. Among the statuses listed, tick **OK**. *OK* is now displayed in the field. A new line appears.
 - f. Click on **APPLY**. The list is now filtered and only the zones that have a status different from *OK* are displayed.
2. In the menu, select  **Alerts, gadgets & Smart Folders > Add an Alert**. The wizard **Add an alert definition** opens.
3. In the field **Name**, name the alert. By default, the alert is named after the module and page from where you configure it, in our example *DNS: Zones*.
4. In the field **Description**, you can type in a description if needed.
5. In the section **Expert mode**, tick the box to display the expert configuration fields.
6. Through the fields **Filter results** and **Value**, you can configure the alert execution parameters.

Table 91.4. Alert execution configuration fields

Field	Description
Filter results	You can select <i>!= (different from)</i> , <i>> (Greater than)</i> , <i>< (Less than)</i> or <i>== (Equal to)</i> . Any of these conditions affects the number typed in the field Value . By default, <i>!= (different from)</i> is selected.

Field	Description
Value	You can type in a number that corresponds to the threshold of your the filter you set before adding the alert. By default, 0 is displayed.

For instance, if you do not want the alert to be triggered for less than 2 zones with a status different from *OK*, you can select *Greater than* in the drop-down list **Filter results** and 2 in the field **Value**.

7. In the section **Triggered by change**, tick the box if you want your alert to match your filter only by change. In the case of our example, if you do not tick the box and three zones already correspond to the filter (they could be in delayed create, timeout...), the alert is triggered if, at the next check, the zones are still not set to *OK*.
8. In the drop-down list **Alert Priority**, define the alert priority. It can be *Low*, *Normal*, *High*, *Urgent* or *Immediate*.
9. In the drop-down list **Alert Severity**, define the alert severity. You can choose among *Minor*, *Major*, *Crash* and *Block*.
10. In the drop-down list **Alert Group Owner**, select a group of users among the ones you created.
11. Tick the box **Edit scheduling** to display the related fields.

Table 91.5. Alert check scheduling parameters

Field	Description
Day(s) of the week	Select a frequency (over the whole week or for a specific set of days) or a specific day of the week.
Date of the month	Select a specific day of the month or a frequency (every day) for the refresh.
Month	Select a specific month or a frequency (every month) for the refresh.
Hour	Select a frequency (over the whole day or for a limited period of time each day), a set of hours or a specific hour per day for the refresh.
Minute	Select the moment (o'clock, quarter past, half past or quarter to) or the frequency (in minutes) of the refresh.

By default, the check is scheduled every 5 minutes of every hour, day and month.

12. Tick the box **Send mail** to display the related fields.

Table 91.6. Email notification configuration parameters

Field	Description
Mailing lists	Select a group of users among the ones created on the page <i>Groups</i> (Administration module). Make sure that the email address of the users belonging to the selected group is configured, otherwise they can never receive the alert.
Additional Mail	Type in the email address of the recipient of the alert and click on ADD to move it to the <i>Additional Mail List</i> . Repeat these actions from as many recipients as needed.
Additional Mail List	In this list are displayed all the recipients of the alert email.

13. Tick the box **SNMP Trap** to display the related fields.

Table 91.7. SNMP trap configuration parameters

Parameter	Description
SNMP version	Select the version of SNMP, either v2c or v3
SNMP Destination	Type in the IP address of the network management platform.
SNMP Community	Type in the community string that would act as a password to access the SNMP agent.

Parameter	Description
Raised alert SNMP OID	Type in a custom OID to be sent when the alert is raised. You can use and extend the default OID <i>1.3.6.1.4.1.2440.1.6.1.2.0.1</i> .
Released alert SNMP OID	Type in a custom OID to be sent when the alert is released. If this field is empty, no trap is sent when the alert is released.

- Click on to complete the operation. It is now listed in the page **Alerts Definition** and marked as **Released**.

Editing Alerts

The alerts can be edited in two different ways:

- You can edit the alert definition: rename it, change the check frequency, add or remove email recipients, set or remove an SNMP trap...
- You can edit the filters that were set on the page when the alert was created.

Editing an Alert Definition

At any time you can edit the properties of your alert.

To edit an alert definition

- In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
- At the end of the line of the alert name of your choice, click on . The properties page opens.
- In the panel **Main properties**, click on . The wizard **Edit an alert definition** opens.
- Edit the alert according to your needs. For more details, refer to the steps 3 to 12 of the procedure [To add an alert](#).
- Click on to complete the operation. The report opens and closes. The properties page is visible again.

Editing an Alert Filters

In addition to editing an alert definition, you can edit the filter(s) it uses. This allows you to edit the alert triggers if the initial alert settings no longer suit your needs.

To edit the filters of an alert

- In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
- Filter the list if need be.
- In the column **Alert filter**, click on *Edit alert filters* for the alert of your choice. You are redirected to the page where the alert was created.
- Above the menu, a blue banner indicates you are in *Alert edition mode*.
- In the search engine of the filtered column(s), edit or remove the filter(s) according to your needs.

7. Add filters based on new columns if relevant.
8. On the right-end side of the menu, click on . The wizard **Quit editing the alert filters** opens.
9. Tick the box **Save changes before quitting** to save your new filter.

If you do not tick the box, your changes are discarded.

10. Click on to complete the operation. The report opens and closes. The page **Alerts Definition** is visible again.

Enabling or Disabling Alerts

If at some point you want an alert definition to stop raising alert instances, you can disable it.

When you want it to raise alert instances again, re-enable it.

To enable/disable an alert

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
3. Tick the alert(s) of your choice.
4. In the menu, select  **Edit > Enable** or **Disable**. The wizard opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again, the alert definition **State** is marked as *Enabled* or *Disabled*.

From the column **State**, you can also directly click on *Enabled* or *Disabled* to, respectively, disable or enable a definition.

Updating an Alert State

At any time, you might want to check that you did not miss any alert. It is useful if you did not configure an alert to trigger on change or if, on the contrary, you just configured it with a check every 5 minutes.

The option *Force alert update* immediately checks if the selected alert is *Raised* or *Released*.

To force an alert definition check

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
3. Tick the alert(s) that you want to check.
4. In the menu, select  **Tools > Force alert update**. The wizard **Force alerts state update** opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again. To see if the alert is now raised, go to the page *Alerts*.

Dismissing an Alert

Once an alert was raised, you can dismiss it in order to make sure that next time it is raised you actually only see the instances that matter and not old ones. The alert does no longer appear on the page *Alerts* unless you tick the box *Display All Alerts*.

To dismiss an alert

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Alerts**. The page **Alerts** opens.
3. Tick the raised alert(s) that you want to dismiss.
4. In the menu, select **Edit > Dismiss**. The wizard **Dismiss an alert** opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again, the alert is no longer listed.
6. Under the menu, tick the box **Display All Alerts** to display all the previous instances.

Reinstating a Dismissed Alert

Once an alert was raised, you can dismiss it as described in the section above. You can however reinstate it as monitored to display it again on the page.

To monitor an alert

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Alerts**. The page **Alerts** opens.
3. Tick the raised alert(s) that you want to reinstate.
4. In the menu, select **Edit > Monitor**. The wizard **Monitor an Alert** opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again, the alert is displayed even though the box *Display All Alerts* is unticked.

Deleting an Alert

For safety measures, you cannot delete an alert instance on its own. However, from the page *Alerts Definition*, you can delete an alert completely i.e. delete its configuration details and the instances of when it was raised.

To delete an alert

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Alerts*, click on **Definition**. The page **Alerts Definition** opens.
3. Tick the alert(s) that you want to delete.
4. In the menu, select **Edit > Delete an Alert**. The wizard **Delete an alert** opens.
5. Click on to complete the operation. The report opens and closes. The alert is no longer listed on the pages *Alerts Definition* and *Alerts*.

Managing the Logs

In the module Administration, two pages allow to manage the logs. You can monitor them from [Syslog](#) and redirect them from [Configuration of Network Logs](#).

Syslog

The page *Syslog* lists the logs of all the services executed. You can filter the list using the menu or the columns to display a specific operation.

Note that you can display the logs of remote appliances from the management SOLIDserver. For more details regarding remote management, refer to the chapter [Centralized Management](#).

To display the logs of your choice on the page Syslog

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
3. Under the menu:

- a. The drop-down list **SOLIDserver** allows to select the SOLIDserver of your choice.

If you are not managing any remote appliance, the list only displays *local*.

If you are managing remote appliances, you can select the local appliance - listed using its hostname - or a remote SOLIDserver - listed using its hostname and IP address as such: *hostname (IP address)*.

- b. The drop-down list **Services** allows to select the service of your choice:

Service	Description
named	The DNS log messages.
dhcpcd	The DHCP log messages.
ipmserver	The internal transactional engine log messages.
messages	All the system log messages.
auth	The authentication log messages.
ipmserver-rules	The operations executed by rules.

- c. The box **Automatic refresh** allows to automate the refresh of all the logs.

By default, the refresh is scheduled to be executed every 10 seconds. To change the refresh frequency, refer to the procedure [To change the automatic refresh frequency](#).

4. Two columns allow to narrow down the log search:

Column	Description
Time	Allows to sort and filter the logs based on the date and time of the service execution ^a .
Log	Allows to filter the logs based on their description.

^aYou can edit the time and date formats via the top bar menu *My Account > My Settings*.

At any time, you can change the Automatic refresh frequency from the registry database.

To change the automatic refresh frequency

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the filter **Name**, type in the *syslog.refresh*. The list is filtered and the registry key *www.system.syslog.refresh* is listed.
4. In the column **Value** of that key, click on the value listed. The default value is *10*. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, replace the current value with the value of your choice (in seconds).
6. Click on to complete the operation. The report opens and closes. The list is visible again and now the automatic refresh occurs at the frequency you just configured.

Configuration of Network Logs

The page *Configuration of Network Logs* allows users of the group *admin* to redirect the logs of several appliances toward a remote syslog server to monitor them.

You can redirect the logs of a particular service and severity level. The available severity levels are listed below.

Table 91.8. Syslog severity levels

Code	Severity level	Description
0 (maximum severity level)	Emergency	The system has completely crashed and is no longer functioning.
1	Alert	The system is unstable and a crash is imminent. Action must be taken immediately.
2	Critical	Critical conditions. Should be corrected immediately.
3	Error	Error conditions. Non-urgent failures that should be relayed to administrators.
4	Warning	Warning conditions. Indicates that an error is returned if no action is taken.
5	Notice	Unusual situation or significant event that is typically part of normal day-to-day operations.
6	Information	Normal operational messages - may be harvested for reporting, measuring throughput, etc - no action required.
7 (minimum severity level)	* (Debug)	Useful messages to developers for debugging, not useful during operations.

Note that selecting a log level automatically includes the logs with a higher severity, the ones with a smaller code number. Therefore, if you select *Warning* (4) logs, you also redirect the *Error* (3), *Critical* (2), *Alert* (1) and *Emergency* (0) logs.

To add a syslog redirection

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Syslog*, click on **Configuration**. The page **Configuration of network logs** opens.

3. In the menu, click on **+** **Add**. The wizard **Syslog configuration** opens.
4. In the drop-down list **Services**, select the service of your choice.

Table 91.9. The services that can be redirected

Service	Description
ipmsvr	The internal transactional engine log messages.
dhcp	The log messages of the service <i>dhcpcd</i> .
dns	The log messages of the service <i>named</i> .
postgres	The SQL log messages.
messages	All the system log messages.
auth	The authentication log messages.
security	The security log messages.
ipmsvr-rules	The operations executed by rules.

5. In the drop-down list **Priority**, select the severity level of your choice. Note that any severity other than *Emergency* (0) also redirects higher severity levels, the ones with a lower code. For more details, refer to the table [Syslog severity levels](#).
6. In the field **Target server**, specify the IP address and port number of the Syslog server receiving the logs following the format `<ip-address>:<port-number>`.
7. Click on **OK** to complete the operation. The report opens and closes. The page **Configuration of Network Logs** is visible again and displays the list of logs redirections.

To delete a syslog redirection

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, next to *Syslog*, click on **Configuration**. The page **Configuration of network logs** opens.
3. Next to the redirection of your choice, click on **DELETE**. The wizard **Syslog configuration** opens.
4. In the fields **Service** and **Target server** are displayed the redirection details. Make sure they are the ones you want to delete.
5. Click on **OK** to complete the operation. The report opens and closes. On the page **Configuration of Network Logs**, the redirection is no longer listed.

Monitoring the Appliance Statistics

The page *System statistics* provides charts dedicated to the appliance services traffic and state. Keep in mind that:

- The system stores data during a year.
- The page allows to display the statistics of your local appliance and the statistics of remote appliances, if they are in version 6.0.1 or higher. For more details regarding remote management, refer to the chapter [Centralized Management](#).
- Every chart displaying local data is a gadget in essence and can be displayed on any dashboard using the dedicated pushpin. For more details, refer to the chapter [Managing Gadgets](#).

- You cannot use the statistics charts of a remote appliance as a gadget.
- You can export all the charts on the page, whether they display local or remote data, in specific reports. For more details, refer to the section [Appliance Statistics Reports](#) below.
- Except for the panel *Processes states*, on each chart you can zoom in and out of the charts or decide the period and data to display. For more details refer to the section [Charts](#).

To access the page dedicated to SOLIDserver Statistics

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **System statistics**. The page **System statistics** opens.

By default, it displays the local appliance data. To display the statistics of a remote appliance, refer to the section [Monitoring the Statistics of Another Appliance](#).

Each chart on the page has a specific purpose:

Table 91.10. The charts of the page Statistics

Chart ^a	Description	Unit of Measurement
DNS traffic	The rate of DNS requests sent and received.	Bytes per second
DHCP traffic	The rate of DHCP requests sent and received.	Bytes per second
HTTP traffic	The rate of HTTP requests sent and received.	Bytes per second
SNMP traffic	The rate of SNMP requests sent and received.	Bytes per second
Database replication traffic	The input and output exchanges during the database replication between two appliances set in a High Availability configuration.	Bytes per second
Load average	The load of all the system's CPUs on an average of 1, 5 and 15 min.	Bytes
CPU per process	The percentage of CPU used by each enabled service.	Percent
Memory usage per process	The memory usage of each enabled service.	Bytes
I/Os per process	The total data input and output of each process.	IO per second
SQL queries	The number of SQL queries made by the system.	Queries per second
Threads	The number of threads executed by the system.	Threads per second
User sessions	The number of users connected at any time.	Connections count
Disk usage	The disk usage.	Percent
Processes state	The state of all the services embedded into SOLIDserver: either <i>OK</i> or <i>down</i> . The DNS and DHCP analytics are part of the processes as well.	Checked every minute

^aThe name and IP address of the appliance can be displayed next to the chart name.

Generating Statistics Reports

EfficientIP provides statistics dedicated reports. The reports on inconsistencies might be empty if the devices configuration is correct.

Table 91.11. Available Statistics reports

Page	Report
System statistics ^a	Statistics chart

Page	Report
	Network traffic

^aOnly users of the group *admin* can access this page.

For more details regarding the reports and their generation, refer to the section [Managing Reports](#).

Tracking Sessions

The page *Session tracking* allows to display the list of the users who recently connected or are currently connected to SOLIDserver. The user connection is checked every 300 seconds.

To track the latest user sessions

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Session tracking**. The page **Session tracking** opens.

You can also track previous sessions on the page *Session history*.

To display the session history of all users

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Session tracking**. The page **Session tracking** opens.
3. On the right-end side of the menu, click on **Session history**. The page **Session history** opens.
4. To display the latest user sessions again, in the breadcrumb click on *Session tracking*.

Tracking Users

The page *User tracking* provides a list of all the operations carried out by every user. The different columns and filters on the page allow to track operations and who performed them. For more details, refer to the sections [Tracking User Operations](#).

You can also grant full access to the page to any group of users. For more details, refer to the section [Allowing Users to Display All the Operations Performed](#).

In addition, note that a registry database entry allows to save the operations performed by users in the file *ipmserver.log*, available on the page Syslog. For more details, refer to the section [Sending a Copy of User Operations to Syslog](#).

Tracking User Operations

To filter the operations and display a specific object addition, deletion or edition you can:

- Use the drop-down list *Services* to select a specific module, object or operation performed. All the available services are listed in the appendix [User Tracking Services Filter](#).
- Use the search engine of the column *Date*, *Service*, *User* and/or *Description*.

To track user operations

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **User tracking**. The page **User tracking** opens.
3. In the drop down list **Services**, filter the results by service if need be. The page refreshes.
4. Filter the list using the search engine of your choice:
 - a. In the column **Date**, you can type in specific date and time or use the filter constructor.
 - b. In the column **Service**, you can filter the operations carried out and described as follows: *<action>:<resource it was performed on>*.
 - c. In the column **User**, you can type in the user who performed the operation.
 - d. In the column **Description**: you can type in details about the object the action was performed on. For the objects with class parameters, you can find the class parameters details. To display all the class parameters details, hover over *Class Parameters*, a pop-up window displays them; or click on *Class Parameters*, all the class parameters details are displayed in the description field along with all the other object details.
5. Click on **REFRESH** to display the corresponding user(s).

Allowing Users to Display All the Operations Performed

By default, any user can access the page User Tracking and see all the changes they performed.

You can grant user access to see the changes performed by all the users, including *ipmadmin*, if their group of users has the permission *User Tracking Display: changes from all the users*.

To grant access to all the changes performed on the appliance to a group of users

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the sidebar, go to **Users, Groups & Rights > Groups**. The page **Groups** opens.
3. At the end of the line of the group of your choice, click on . The properties page opens.
4. In the panel **Administration**, click on **EDIT**. The wizard **Edit group rights** opens.
5. The **Unauthorized services** list displays the services that are not granted to the group. Select *User Tracking Display: changes from all the users* and click on . The service is moved to the list **Authorized services**.
6. Click on **OK** to complete the operation. The report opens and closes. The page refreshes. In the panel, the list **Permissions** displays the service.

Once the permission is granted, all the users of the group can see the operations performed by anyone who logged in SOLIDserver on the page *User tracking* and in the panel *Audit* of the properties page of an object.

Sending a Copy of User Operations to Syslog

SOLIDserver allows you to save a copy of users activity in the file *ipmserver.log*, using a dedicated registry key. This way, user operations are available both on the pages Syslog and User Tracking.

To send a copy of user operations to Syslog

1. Add the registry key to enable the external storage of user operations

Only users of the group *admin* can perform this operation.

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- c. In the column **Name**, type in *module.system.usertracking_use_syslog* and hit Enter. The key is the only one listed.
- d. In the column **Value**, click on *0*. The wizard **Registry database Edit a value** opens.
- e. In the field **Name**, the key name is displayed in a read-only gray field.
- f. In the field **Value**, delete the *0* and replace it with a *1*. This value means the key is enabled.
- g. Click on to complete the operation. The report opens and closes. In the column **Value**, a *1* is displayed.

2. Display the user operations in Syslog

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- b. In the section **Monitoring**, click on **Syslog**. The page **Syslog** opens.
- c. In the drop-down list **SOLIDserver**, verify that the local appliance is selected. Only the hostname appears with no IP address.
- d. In the drop-down list **Services**, select *ipmsserver*. The page refreshes.
- e. In the column **Log**, use the filter *ipmsserver:* . The user operations are listed as follows:

```
<hostname> <process_name>[<process_id>]: ipmsserver: <service_name> <user_name>
<service_parameters>
```

Managing SNMP Profiles

SNMP profiles are used to collect SNMP data from hosts or other devices running an SNMP or proxy SNMP agent. SNMP profiles allow you to monitor remotely DHCP and DNS services through the SNMP protocols. For more details, refer to the section [Managing the SNMP Service](#).

By default, SOLIDserver already provides 3 SNMP profiles (*standard v1*, *standard v2c* and *standard v3*). To edit these profiles, refer to the section [Editing an SNMP Profile](#).

Adding an SNMP Profile

To add an SNMP profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.

3. In the panel **SNMP profiles configuration**, click on **ADD**. The wizard **Add an SNMP profile** opens.
4. In the field **SNMP profile name**, name the profile.
5. In the field **Description**, you can type in a description.
6. In the drop-down list **SNMP version**, select the SNMP version you want to use.
7. Click on **NEXT**. The next page opens.
8. If you selected the *v1* or *v2c* version of SNMP:
 - a. In the field **Read community**, type in the read-only community string that would act as a password for this profile reading requests. For the preexisting profiles *standard v1* and *standard v2*, the default value is *public*.
 - b. In the field **Write community**, you can type in a write community string that would act as a password for this profile reading and writing requests. For the preexisting profiles *standard v1* and *standard v2*, the default value is *private*.
9. If you selected the *v3* version of SNMP, fill in the **Read access parameters** and the **Write access parameters** according to the table below:

Table 91.12. SNMP v3 profiles access parameters

Field	Description
User name	In this field, type in the user name. This field is required for read access parameters. For the preexisting profile <i>standard v3</i> , the default value is <i>default_ipm_user</i> .
Authentication key	In this field, type in a key to ensure the authentication of the source. This field is required for read access parameters. For the preexisting profile <i>standard v3</i> , the default value is <i>default_auth_key</i> .
Authentication	In this field, select the cryptographic hash function used for authentication: either <i>MD5</i> , <i>SHA</i> or <i>None</i> . This field is required for read access parameters. For the preexisting profile <i>standard v3</i> , the default value is <i>MD5</i> .
Privacy key	In this field, if need be, type in the encryption key to prevent snooping from unauthorized sources.
Privacy	In this field, if need be, select the encryption type: either <i>DES</i> or <i>None</i> . For the preexisting profile <i>standard v3</i> , the default value is <i>DES</i> .

10. Click on **OK** to complete the operation. The page **SNMP profiles configuration** is visible again, your profile is listed in the panel.

Editing an SNMP Profile

To edit an SNMP profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.
3. In the panel **SNMP profiles configuration**, select the *SNMP profile configuration* you want to edit.
4. Click on **EDIT**. The **Edit an SNMP profile** wizard opens.
5. In the field **Description**, you can type in a description.

6. In the drop-down list **SNMP version**, select the SNMP version you want to use.
7. Click on **NEXT**. The next page opens.
8. If you are editing a profile in SNMP v1 or v2c: edit the **Read community** and/or field **Write community** as needed.
9. If you are editing a profile in SNMP v3, edit the **Read access parameters** and **Write access parameters** as needed.
10. Click on **OK** to complete the operation. The page **SNMP profiles configuration** is visible again.

Deleting an SNMP Profile

To delete an SNMP profile

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Network devices & SNMP profiles**. The page **Network devices & SNMP profiles** opens.
3. In the panel **SNMP profiles configuration**, select the *SNMP profile configuration* you want to delete.
4. Click on **DELETE**. The wizard opens.
5. Click on **OK** to complete the operation. The page **SNMP profiles configuration** is visible again. The profile has been removed from the **SNMP profiles configuration** list.

Monitoring Using SNMP

SNMP monitoring allows to supervise SOLIDserver components and processes to ensure they meet the expected performance and availability.

Combined with an effective alerting system, it allows to gather and store metrics about key system health indicators to analyze and correlate information. Based on that data, you can intervene in case of malfunction, define alert triggers and set up automatic actions that prevent an overall failure of the system.

You can monitor these metrics through an external solution such as Nagios⁵ as well as some related plug-ins.

Keep in mind that:

- You will need an Internet connection and your credentials to access the files *.mib* on our website.
- Each monitored SOLIDserver should be configured to allow the SNMP collector, which must use SNMP v2c or v3, to retrieve the SNMP information.

Once your system is properly configured, you can set various SNMP alerts on SOLIDserver objects to be notified of any unusual behavior. For more details, refer to the chapter [Managing Alerts](#).

For more details about the available metrics, refer to the appendix [SNMP Metrics](#).

⁵<https://www.nagios.org>

Displaying Netstat Data

SOLIDserver provides a page listing Netstat data. This tool allows to display the open TCP and UDP ports to monitor active connections on the management appliance.

To access the page Netstat

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Monitoring**, click on **Netstat**. The page **Netstat** opens.

This page contains several columns:

Table 91.13. Netstat columns

Column	Description
Protocol	The protocol name, TCP or UDP.
Local address	The local appliance IP addresses. That is to say any IP addresses configured for your physical interfaces - for more details, refer to the chapter Configuring the Network - and the loopback IP address.
Local port	The number of the local appliance port through which the connection is made. If the port is not yet established, an asterisk (*) is displayed.
Foreign address	The IP address and port number of the remote computer to which the socket is connected.
Foreign port	The number of the remote appliance port through which the connection is made. If the port is not yet established, an asterisk (*) is displayed.
State	The state of the TCP connection. It can be: <ul style="list-style-type: none"> <i>LISTEN</i>. The socket is listening for incoming connections. <i>ESTABLISHED</i>. The socket has an established connection. <i>SYN_SENT</i>. The socket is actively trying to establish a connection. <i>TIME_WAIT</i>. The socket is waiting after close to handle packets still in the network.

Sizing the Database Tables

SOLIDserver provides a page that lists the size of all the tables in the database. This list gives you all the information in one glance as it includes the **Table name**, **Total size** (including the index), **Table size** and **Tuple size** columns to even provide you with the size of the data and tuples they contain.

To access the page Database tables size

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Database tables size**. The page **Database tables size** opens.

Vacuuming the Database

By default, SOLIDserver enables the rule 180 that vacuums the database every night at 4 a.m. This rule is configured to defragment the databases and reclaim storage occupied by dead tuples to increase performances of the processes related to users queries.

Users of the group *admin* can edit the rule to change the database vacuuming frequency. **You should only edit this rule if a member of an EfficientIP technical team specifically asked for it.**

To edit the rule 180 that defragments SOLIDserver database

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Rules**. The page **Rules** opens.
3. In the column **Rule #**, type in *180* and hit Enter. The rule is the only one listed.
4. At the end of the rule line, click on . The properties page opens.
5. In the panel **Main properties**, click on **EDIT**.
6. In the fields **Module**, **Event**, **Rule** and **Description** are details the rule properties.
7. In the field **Rule name**, you can edit the name of the rule. It is displayed in the column *Instance* on the page *Rules*.
8. In the field **Comment**, you can type in a comment if you want.
9. Click on **NEXT**. The page **Rule filters** opens.

Set the schedule parameters:

Table 91.14. Scheduled rules parameters

Field	Description
Day(s) of the week	Select a day or a period of days. By default, <i>Every day</i> is selected.
Date of the month	Select a date. By default, <i>Every day</i> is selected.
Month	Select a month. By default, <i>Every month</i> is selected.
Hour	Select a specific time or one of the available schedules. By default, <i>04</i> is selected.
Minute	Select a period of time, minutes-wise. By default, <i>00</i> is selected.

10. Click on **NEXT**. The page **Rule parameters** opens.
11. In the field **Max. wasted space (MB)**, type in the maximum size of wasted space beyond which a database table is vacuumed. By default, it is set to *200*.
12. Click on **OK** to complete the operation. The report opens and closes. The changes are displayed.

Chapter 92. Maintenance

SOLIDserver needs to be properly maintained over time to run smoothly and reach its maximum performance. This chapter details the pages and options that allow administrators to manage certificates and local files, troubleshoot SOLIDserver or even enable the Maintenance mode.

This chapter gathers the following maintenance possibilities:

- [Managing the HTTPS Certificate](#).
- [Managing Files from the Local Files Listing](#).
- [Using the Maintenance mode](#).
- [Updating the Macros and Rules](#).
- [Clearing the Appliance Cache](#).
- [Troubleshooting](#).
- [Managing Backups and Restoring Configurations](#).
- [Shutting Down and Rebooting](#).

Managing the HTTPS Certificate

From the page *All Certificates*, you can add X.509 certificates (official and auto-signed SSL certificates), private keys and CSR files (certificate signing requests). You can [import SSL files](#), create [self-signed certificates](#), [CSR files](#) or [private keys](#) and [download](#) the certificates, public and private keys associated with them.

As the self-signed certificate used by default on the page *Services configuration* is not trusted by your web browser, warning messages appear to inform you that the certificate is not from a trusted certifying authority, that the hostname of the certificate is invalid, etc. This connection can be prone to a man-in-the-middle (MITM) attack.

When you receive such warnings, you can accept the certificate just for the current session and save it in the certificate store of your browser.

You can also authenticate SOLIDserver and eliminate the certificate warnings altogether if you:

- Create a self-signed certificate through the GUI or import a CA signed certificate as detailed in the sections [Creating Self-signed Certificates](#) or [Importing an SSL Certificate](#) below.
- Change the certificate. For more details, refer to the sections [Configuring the HTTPS Certificate](#) in the chapter *Configuring the Services*.

Note that **the SSL certificate is unique to each SOLIDserver appliance**. So if you want to use a certificate and you are managing remote appliances or appliances in High Availability: use the drop-down list *SOLIDserver* to make sure you are setting each appliance with its own SSL certificate.

Importing SSL Objects

You can import self-signed or CA signed certificates, CSRs or private keys.

For instance, once the CA sent you a certificate, you can import it to SOLIDserver as described below.

Prerequisites

- You can only import files that do not include any passphrase. If they do, the HTTP protocol cannot start and you might lose the GUI access to your appliance.
- Self-signed and CA signed certificates must be imported via a *.TAR archive file containing:
 - the certificate file respecting the CRT format and named *certificate* without extension.
 - the private key file named *private_key* without extension.
- CSR files must be imported via a *.TAR archive file containing:
 - the request file named *csr_file* without extension.
 - the private key file named *private_key* without extension.
- Private keys must be imported via a file named *object_key* without extension.

To import an SSL object

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
3. In the menu, select  **Import** > **Import**. The wizard **Import an SSL certificate** opens.
4. In the drop-down list **Object type**, select required type.
5. Click on and look for the certificate. Once selected, it is visible in the field **File name**.
6. Click on to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

If you imported a certificate (CA signed or self-signed), you can choose to use it as HTTPS certificate on the page *Services configuration*. For more details, refer to the section [Configuring the HTTPS Certificate](#).

If you imported a private key, you can use it to create a certificate or a CSR. For more details, refer to the sections [Creating Self-signed Certificates Using an Existing Private Key](#) or [Creating CSRs Using an Existing Private Key](#).

Creating Self-signed Certificates

You can create a self-signed certificate automatically generating a private key or using an existing private key.

Creating Self-signed Certificates automatically Generating a Private Key

You can create self-signed certificates respecting the X509 format. During the addition, a private key is automatically generated.

To create a self-signed certificate automatically generating a private key

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.

3. In the menu, click on **+** **Add**. The wizard **Create an SSL object** opens.
4. In the field **Object Name**, name your certificate.
5. In the drop-down list **SSL File Type**, select *X509 certificate*.
6. In the drop-down list **Encryption type**, *RSA* is displayed in read-only.
7. In the field **Encryption**, type in the value of your choice. By default, *2048* is displayed.
8. In the field **Certificate Validity (days)**, edit the number of days if need be. By default, *1825* is displayed.
9. In the drop-down list **Digest method**, select *SHA224*, *SHA256*, *SHA384*, *SHA512* or *DSS1*.
10. Click on **NEXT**. The last page of the wizard opens.
11. Configure the file details:
 - a. In the field **Country Code**, type in the two letter code of your country.
 - b. In the field **State or Province**, type in the state, province or region name in full letters.
 - c. In the field **Locality**, type in the city name.
 - d. In the field **Organization Name**, type in your company name.
 - e. In the field **Organization Unit Name**, type in the name of the department final user among the company.
 - f. In the field **Common Name**, type in the appliance hostname.
 - g. In the field **Email address**, type in your email address.
12. Click on **OK** to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

Once you have created an X509 certificate, you can use it as HTTPS certificate. For more details, refer to the section [Configuring the HTTPS Certificate](#).

Creating Self-signed Certificates Using an Existing Private Key

You can create self-signed certificates using private keys that were automatically generated when creating another certificate, that you imported or that you created in the GUI.

For more details about how to create a private key in the GUI, refer to the section [Creating Private Keys](#).

To create a self-signed certificate using an existing private key

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
3. In the menu, click on **+** **Add**. The wizard **Create an SSL object** opens.
4. In the field **Object Name**, name your certificate.
5. In the drop-down list **SSL File Type**, select *X509 certificate*.
6. Tick the box **Use a previously generated private key**. The field **Use key** appears.

7. In the drop-down list **Use key**, select one of your private keys or certificates or even the default entry *Apache SSL Key Base* or *request_auto_key*.
8. In the field **Certificate Validity (days)**, edit the number of days if need be. By default, *1825* is displayed.
9. In the drop-down list **Digest method**, select *MD5*, *SHA1* or *MD2*.
10. Click on **NEXT**. The last page of the wizard opens.
11. Configure the file details:
 - a. In the field **Country Code**, type in the two letter code of your country.
 - b. In the field **State or Province**, type in the state, province or region name in full letters.
 - c. In the field **Locality**, type in the city name.
 - d. In the field **Organization Name**, type in your company name.
 - e. In the field **Organization Unit Name**, type in the name of the department final user among the company.
 - f. In the field **Common Name**, type in the appliance hostname.
 - g. In the field **Email address**, type in your email address.
12. Click on **OK** to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

Once you have created an X509 certificate, you can use it to change the HTTPS certificate as detailed in the section [Configuring the HTTPS Certificate](#).

Creating CSRs

You can create CSR files automatically generating a private key or using an existing private key to request a certificate respecting the CRT format.

Creating CSRs automatically Generating a Private Key

You can create Certificate Signing Requests (CSR), during the addition, a private key is automatically generated.

To create a CSR automatically generating a private key

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
3. In the menu, click on **+ Add**. The wizard **Create an SSL object** opens.
4. In the field **Object Name**, name your certificate.
5. In the drop-down list **SSL File Type**, select *CSR File*.
6. In the drop-down list **Encryption type**, *RSA* is displayed in read-only.
7. In the field **Encryption**, type in the value of your choice. By default, *2048* is displayed.
8. Click on **NEXT**. The last page of the wizard opens.

9. Configure the file details:
 - a. In the field **Country Code**, type in the two letter code of your country.
 - b. In the field **State or Province**, type in the state, province or region name in full letters.
 - c. In the field **Locality**, type in the city name.
 - d. In the field **Organization Name**, type in your company name.
 - e. In the field **Organization Unit Name**, type in the name of the department final user among the company.
 - f. In the field **Common Name**, type in the appliance hostname.
 - g. In the field **Email address**, type in your email address.
10. Click on to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

Once you have created a CSR, you need to download it and send it to the Certificate Authority. For more details, refer to the section [Downloading SSL objects](#).

Creating CSRs Using an Existing Private Key

You can create CSR Files using private keys that were automatically generated when creating a certificate, that you imported or that you created in the GUI.

For more details about how to create a private key in the GUI, refer to the section [Creating Private Keys](#).

To create a CSR using an existing private key

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
3. In the menu, click on  **Add**. The wizard **Create an SSL object** opens.
4. In the field **Object Name**, name your certificate.
5. In the drop-down list **SSL File Type**, select *CSR File*.
6. Tick the box **Use a previously generated private key**. The field **Use key** appears.
7. In the drop-down list **Use key**, select one of your private keys or certificates or even the default entry *Apache SSL Key Base* or *request_auto_key*.
8. Click on . The last page of the wizard opens.
9. Configure the file details:
 - a. In the field **Country Code**, type in the two letter code of your country.
 - b. In the field **State or Province**, type in the state, province or region name in full letters.
 - c. In the field **Locality**, type in the city name.
 - d. In the field **Organization Name**, type in your company name.
 - e. In the field **Organization Unit Name**, type in the name of the department final user among the company.

- f. In the field **Common Name**, type in the appliance hostname.
 - g. In the field **Email address**, type in your email address.
10. Click on to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

Once you have created a CSR, you need to download it and send it to the Certificate Authority. For more details, refer to the section [Downloading SSL objects](#).

Creating Private Keys

In addition to private keys automatically generated when creating certificates or CSR, you can create new ones and use them to create certificates or CSRs as previously described.

For more details, refer to the sections [Creating Self-signed Certificates Using an Existing Private Key](#) and [Creating CSRs Using an Existing Private Key](#).

To create a private key

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens.
3. In the menu, click on  **Add**. The wizard **Create an SSL object** opens.
4. In the field **Object Name**, name your certificate.
5. In the drop-down list **SSL File Type**, select *Private Key*.
6. In the drop-down list **Encryption type**, *RSA* is displayed in read-only.
7. In the field **Encryption**, type in the value of your choice. By default, *2048* is displayed.
8. Click on to complete the operation. The report opens and closes. The certificate is listed on the page *All certificates*.

Downloading SSL Objects

From the properties page of SSL Objects, you can download certificates, private keys and/or public keys.

For certificates and CSR files, the panels *Certificate*, *Private key* and *Public key* are available. For private keys, the panels *Certificate* and *Private key* are available.

To download certificates, private keys or public keys

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Authentication & Security**, click on **Certificates and keys**. The page **All certificates** opens. The *Apache SSL Cert Base* is listed.
3. At the end of the line of the file of your choice, click on . The properties page opens.
4. In the panel *Certificate*, *Private Key* or *Public Key*, click on and save the file.

Once you have downloaded the certificate associated to a CSR, you can send it to the CA and then import the CA-signed certificate. For more details, refer to the section [Importing an SSL Certificate](#).

Managing Files from the Local Files Listing

The page **Local files listing** allows users of the group *admin* to manage all types of files uploaded or stored locally on the appliance.

All the files are separated among 6 subpages: *Local*, *TFTP*, *Logs*, *Config files*, *Custom images* and *Custom WSDL*. From each of these pages, you can upload, download and delete local files. For more details, refer to the section [Managing Local Files](#) below.

Browsing the Local Files Listing Database

From the page *Local Files Listing*, you can display any local files subpage that suits your needs.

To display the pages of the Local Files listing

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Under the menu, tick the radio button **Local**, **TFTP**, **Logs**, **Config files**, **Custom images** or **Custom WSDL**. The page opens. By default, the list *Local* is displayed.

All subpages, except *Custom WSDL*, share a common set of columns but contains specific files.

Table 92.1. Columns on the page Local files listing

Column	Description
Name	The entry name and extension if relevant. It is underlined as you can display a directory content or download a file.
Type	Indicates if the entry listed in a <i>File</i> or <i>Directory</i> .
Mode	The entry permissions.
Owner	The entry owner i.e. the user logged when the entry was generated.
Group	The file or directory group.
Size	The file or directory size in <i>B</i> , <i>kB</i> or <i>MB</i> .
Last Modified	The month, day, date and time of the last update of the entry or its upload date.

The Page Local

It displays all the files stored locally in the appliance, including:

- The **scheduled exports** configured from any page of the GUI where the menu *Report > Export* is available. Their extension depends on the chosen export file: *.csv*, *.html*, *.xml*, *.xls* or *.pdf*. For more details, refer to the chapter [Exporting Data](#).
- The **reports** generated from the GUI. Their extension depend on the chosen file format: either *.html* or *.pdf*.
- The file **sysaudit.log** that stores in real time all the appliance system information (memory use, partition, netstats, etc). To download this file, refer to the section [Downloading Files](#) below.

- The **network devices captures**. The captures extension is .pcap. For more details, refer to the section [Making a Network Device Snapshot](#).
- The **corrupted configuration files** that triggered a Locked synchronization. For more details, refer to the DNS section [Handling the Status Locked Synchronization](#) or the DHCP section [Handling the Status Locked Synchronization](#).
- The **troubleshooting dump files** generated from the Administration homepage. For more details, refer to the section [Troubleshooting Dump](#).

The Page TFTP

This list displays all the files uploaded locally, available for download, and the files uploaded remotely via TFTP. For more details, refer to the section [Managing the TFTP Upload Authorizations](#).

The Page Logs

It displays all the appliance log files in alphabetical order. To browse their content, go to the page [Syslog](#). For more details, refer to the section [Syslog](#).

The Page Config files

It displays all the servers configuration files generated from the Services configuration page. For more details, refer to the section [Downloading the DNS/DHCP/DHCPv6 Configuration File](#).

The Page Custom images

It displays all the images that you uploaded to customize SOLIDserver login page and home page Welcome banner. For more details, refer to the section [Uploading an Image to SOLIDserver](#).

The Page Custom WSDL

It displays all the WSDL files available for web services management via SOAP on the appliance. The page contains a specific set of columns:

Table 92.2. Custom WSDL Columns

Column	Description
Name	The entry name. Clicking on the name opens the list of services configured in the file.
Endpoint	The chosen Endpoint set for the WSDL file.
Creation date	The file creation time and date.
Dump date	The file dump time and date.
Status	The file status: <i>Dumped</i> , <i>Modified since dump</i> , <i>IP doesn't exist</i> or <i>Dumping in progress</i> .

Clicking on a WSDL file allows to display all the services it contains. For more details, refer to the section [Managing WSDL Files](#).

Managing Local Files

From all the pages of the Local Files Listing you can download and delete files. You can only upload files from the pages *Local*, *TFTP* and *Custom images*.

The page *Custom WSDL* contains a set of specific options detailed in the section [Managing WSDL Files](#).

Uploading Files

From the pages *Local*, *TFTP* and *Custom images* you can upload files. This upload updates the appliance local database from the GUI.

To upload a file to the page **Local files listing**

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Under the menu, tick the radio button *Local*, *TFTP* or *Custom images*. The page opens.
4. In the menu, select . **Tools > Upload file**. The wizard **Import a file** opens.
5. Click on to select the file to upload from your local file system.
6. Click on to complete the operation. The report opens and closes. The file can be downloaded from the page **Local files listing** available from the page *Admin Home*.

Downloading Files

Any file listed on the Local Files Listing can be downloaded to your local computer from the GUI.

To download a file from the page **Local files listing**

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Under the menu, tick the radio button of your choice. The page opens.
4. In the column **Name**, filter the list if need be.
5. Click on the name of the file of your choice to download it¹.

Deleting Files

From any page of the Local files Listing you can delete files from the appliance local database.

To delete a file from the page **Local files listing**

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Under the menu, tick the radio button of your choice. The page opens.
4. Filter the list if need be.
5. Tick the file(s) you want to delete.
6. In the menu, select . **Edit > Delete file(s)**. The wizard **Delete file** opens.

¹Depending on your browser, you might download the file right away or be offered the possibility to open the file or save it.

7. Click on **OK** to complete the operation. The report opens and closes. The file is no longer listed.

Managing WSDL Files

The Web Services Description Language files are used to call services with SOAP.

By default a *full* WSDL file that includes all the services is available (*https://<ip_address>/interfaces/wsd_eip_full.wsd*). Keep in mind that using it drastically reduces the performances of the application you develop to execute the services. Therefore, we recommend creating your own files to only contain the information relevant to the required services.

You can create as many WSDL files as you need, they are locally saved in the directory *https://<SOLIDserver-ip-address>/interfaces/custom_wsd/*. Once created, you must include the files location in your source code to call the services they contain, and preferably make sure to copy and store your custom files locally on the SOAP client to reduce the latency of instantiating the SOAP object

You can add, edit, dump and delete WSDL files from the page *Custom WSDL* of the Local files listing.

Adding a Custom WSDL File

You can create WSDL files that contain as many services as you need. The addition wizard provides an auto-completion field that recognizes the services' name. Once created:

- WSDL files are automatically dumped, you can call the services listed in the WSDL file no matter what programming language you use.
- WSDL files are saved locally in the directory *custom_wsd* in the format *<WSDL-file-name>.wsdl*.

To add a custom WSDL file through SOLIDserver GUI

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. In the menu, click on **+ Add**. The wizard **Add a custom WSDL file** opens.
5. In the field **WSDL file name**, name the file without extension, it is automatically in the format *<file-name>.wsdl*.
6. In the drop-down list **WSDL EndPoint**, select the EndPoint of your choice. It can be your SOLIDserver *Hostname* or IP address.
7. In the field **Service name**, type in the first letters of a service. The auto-completion provides a list of all the services matching what you specified.
8. Select the service of your choice and click on **ADD**. The service name is moved to the list **Selected services**.
9. Repeat these actions for as many services as you want.
 - To update an entry: select the service in the list **Selected services**, edit the service name in the field **Service name**, and click on **UPDATE**. The new service name replaces the former one.

- To delete an entry: select the service in the list **Selected services** and click on **DELETE**.
 - To discard the latest changes, click on **CANCEL**.
10. Click on **OK** to complete the operation. The report opens and closes. The file is listed on the page and saved in the directory *custom_wsdl*. To display the list of services it contains refer to the next procedure.
 11. To use the services listed in the WSDL file, you must integrate the file location to your source code using its absolute address within SOLIDserver database: *https://<ip_address>/interfaces/custom_wsdl/<your-WSDL-file-name>.wsdl*.

The services that you configured in the file are all listed on a dedicated page.

To display the content of a WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. Click on the name of the WSDL file of your choice. The page **Content of the WSDL file: <selected-file>** appears and displays all the services that the file contains in the column **Service name**.
5. To go back to the page Custom WSDL tick the radio button **Custom WSDL**. The page opens.

Editing and Dumping a Custom WSDL File

At any point you can edit the content of a WSDL file, that is to say add or remove services. Keep in mind that **editing a WSDL file requires dumping it again**.

Adding Services to an Existing Custom WSDL File

You can add as many services as you want to an existing custom WSDL file. After adding services, you must dump the WSDL file to take into account the changes.

To add services to a custom WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. Click on the name of the WSDL file of your choice. The page **Content of the WSDL file: <selected-file>** appears and displays all the services that the file contains in the column **Service name**.
5. In the menu, click on **+ Add**. The wizard **Add a service** opens.
6. In the field **Service name**, type in the first letters of a service. The auto-completion provides a list of all the services matching what you specified.
7. Select the service of your choice and click on **ADD**. The service name is moved to the list **Selected services**. Repeat these actions for as many services as you want.

- To update an entry: select the service in the list **Selected services**, edit the service name in the field **Service name**, and click on **[UPDATE]**. The new service name replaces the former one.
 - To delete an entry: select the service in the list **Selected services** and click on **[DELETE]**.
 - To discard the latest changes, click on **[CANCEL]**.
8. If you want to dump the WSDL file upon edition:
 - a. Tick the box **Dump WSDL file(s)**. The field *WSDL Endpoint* appears.
 - b. In the drop-down list **WSDL EndPoint**, you can edit the WSDL file EndPoint. By default, the current Endpoint is selected.
 9. Click on **[OK]** to complete the operation. The report opens and closes. The new services are listed.

To go back to the page Custom WSDL tick the radio button **Custom WSDL**. The page opens.

If you ticked the box *Dump WSDL file(s)*, the file's **Status** is *Dumped* so the file is updated and can be used immediately. Otherwise, it is marked *Modified since dump*, and you need to dump the file to take into account your changes, for more details refer to the section [Dumping an Edited Custom WSDL File](#).

Removing Services from an Existing Custom WSDL File

You can remove one or several services from an existing WSDL file. After removing services, you must dump the WSDL file to take into account the changes.

To remove services from a custom WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. Click on the name of the WSDL file of your choice. The page **Content of the WSDL file: <selected-file>** appears and displays all the services that the file contains in the column **Service name**.
5. Tick the service(s) you want to delete.
6. In the menu, select **↗ Edit > Delete a service from a WSDL file**. The wizard **Delete service** opens.
7. If you want to dump the WSDL file upon edition:
 - a. Tick the box **Dump WSDL file(s)**. The field *WSDL Endpoint* appears.
 - b. In the drop-down list **WSDL EndPoint**, you can edit the WSDL file EndPoint. By default, the current Endpoint is selected.
8. Click on **[OK]** to complete the operation. The report opens and closes. The services are no longer listed.

If you ticked the box *Dump WSDL file(s)*, the file's **Status** is *Dumped* so the file is updated and can be used immediately. Otherwise, it is marked *Modified since dump*, and you need to dump the file to take into account your changes, for more details refer to the section [Dumping an Edited Custom WSDL File](#).

To go back to the page Custom WSDL tick the radio button **Custom WSDL**. The page opens.

Dumping an Edited Custom WSDL File

If you edited a WSDL file without dumping it again in the edition wizard, you can dump its content, or change its Endpoint if need be, from the menu on the page *Custom WSDL*.

To dump a custom WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. In the column **Name**, filter the list if need be.
5. Tick the file(s) you want to dump.
6. In the menu, select  **Tools > Dump WSDL File(s)**. The wizard **Dump WSDL file(s)** opens.
7. If you want to change the EndPoint:
 - a. Tick the box **Dump WSDL file(s)**. The field *WSDL Endpoint* appears.
 - b. In the drop-down list **WSDL EndPoint**, you can edit the WSDL file EndPoint. By default, the current Endpoint is selected.
8. Click on to complete the operation. The report opens and closes. The files are marked as *Dumped* in the column **Status**. If you edited the EndPoint, its new value is visible in the column **Endpoint**.

Deleting a Custom WSDL File

At any time, you can delete a custom WSDL file that you no longer use.

To delete a custom WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. Tick the file(s) you want to delete.
5. In the menu, select  **Edit > Delete file(s)**. The wizard **Delete file** opens.
6. Click on to complete the operation. The report opens and closes. The file is no longer listed.

Downloading a Custom WSDL File

At any time, you can download a WSDL file. To download your file, you must dump it on the page *Local* and download it from there.

To download a custom WSDL file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Above the menu, tick the radio button **Custom WSDL**. The page opens.
4. In the column **Name**, filter the list if need be.
5. Tick the file(s) you want to dump.
6. Dump the file on the page *Local*:
 - a. In the menu, select **Tools > Dump in Local Files Listing**. The wizard opens and closes.
 - b. The file is now listed in the page *Local* of the Local Files Listing.
7. Download the file:
 - a. Above the menu, tick the radio button **Local**. The page opens.
 - b. In the list, the WSDL is listed as follows: *<WSDL-file-name>.tar* .
 - c. In the column **Name**, click on the file to download it².

Using the Maintenance mode

The Maintenance mode allows members of the group *admin* to work without interferences on their infrastructures. It disconnects all non-admin users from SOLIDserver during maintenance work, like reorganizing the network infrastructure or modifying services configuration as users intervention may affect the administrators actions.

Keep in mind that:

- Enabling the Maintenance mode does not interrupt network services.
- In Maintenance mode, the users that do not belong to the group *admin* are disconnected and/or cannot log in. Once the mode is disabled, them can connect again.
- Users of the group *admin* that are connected when the mode is enabled have a red banner message is displayed on every page of SOLIDserver to inform them that the mode is on.

To enable/disable the Maintenance mode

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

²Depending on your browser, you might download the file right away or be offered the possibility to open the file or save it.

2. In the section **Maintenance**, click on **Maintenance Mode**. The wizard **Enable/Disable Maintenance mode** opens.
3. Click on to complete the operation. The report opens and closes. The page **Admin Home** is visible again.

Updating the Macros and Rules

When you add new macros and rules, usually for customization purposes, SOLIDserver must take into account the files that describe them. Therefore, users of the group *admin* might have to register the new macros and rules into the system. This operation is usually supervised by the EfficientIP support team.

To update the macros and rules

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Register new macros & rules**. The wizard **Register all the latest macros and rules** opens.
3. Click on to complete the operation. The report opens and closes. The page **Admin Home** is visible again.

Clearing the Appliance Cache

In case of changes in SOLIDserver code, for instance when a hotfix was provided, users of the group *admin* may have to reload the file system cache. This operation is usually supervised by the EfficientIP support team.

To clear SOLIDserver's cache

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Clear SOLIDserver cache**. The wizard **Clear SOLIDserver cache** opens.
3. Click on to complete the operation. The report opens and closes. The page **Admin Home** is visible again. Any internal modification of the code has been taken into account.

Troubleshooting

Troubleshooting is a logical and systematic search for the source of a problem. It is needed to develop and maintain complex systems where symptoms can have many possible causes.

Before Troubleshooting

There is set of simple checks that might help you avoid troubleshooting. These checks are often overlooked in times of functional problems when they should be an administrator reflex.

1. **Make sure that the appliance and the objects it manages are set at the same time.** If they are not, set the appliance time.

Typically, if your appliances and the servers it manages are not the same time, you can encounter management problems: the DHCP should be the first impacted with the leases and then the DNS, especially if you set time check keys for the zones. We recommend that you configure NTP servers on the appliance as detailed in the section [Configuring NTP Servers](#). Besides, we strongly advise against setting the time through CLI because it might make SOLIDserver crash, disrupt your services, trigger errors in the logs, etc. If you do it anyway, restart SOLIDserver to make sure that all the services impacted by the time change are restarted and all set at the same time.

2. Make sure there is no Multi-Management of your DNS and DHCP physical servers.

Through the smart architectures, you can manage the servers of your choice so make sure you did not add and manage twice the same server in two different smart architectures. Every minute the smart architecture checks that its configuration is pushed to the physical server, if not it pushes it again. So if one physical server is managed through two different architectures every minute a configuration is pushed and then overwritten by the other smart architecture.

Troubleshooting Guidelines

Determining what might be the causes of a dysfunction is often a process of elimination. Troubleshooting also requires confirmation that the solution restores the system to its working state.

The following guidelines give a generic overview of troubleshooting, and since each case is different, you might need to vary your approach to the problem.

How to troubleshoot your system

1. **Confirm the presence of a backup** in case of service interruption. You might need the backup file to restore the previous stable version of your system. However, restoration overwrites the changes made between the time of the backup and the time of the crash, so this would be the very last resort. For more details, refer to the section [Managing Backups and Restoring Configurations](#).
2. **Isolate the malfunctioning behavior** to pinpoint what services or components are affected.
3. **Inspect the status indicators** that can highlight a dysfunction.
4. **Inspect connections** to any attached devices and check their power sources.
5. **Review the network and services configuration**. For more details, refer to the part [Configuring SOLIDserver](#).
6. **Check if the issue is not due to the customer background**, i.e. the customer's use of the services, operating system, network topology components and levels of software that were running when the incident occurred.
7. **Check the product logs**. Do not hesitate to check the DNS logs, DHCP logs, PostgreSQL logs, the management logs as well as the system logs. For more details, refer to the section [Syslog](#).
8. **Check the system logs**. Do not hesitate to check the *sysaudit.log* file, available on the Local Files Listing page. For more details, refer to the section [Managing Files from the Local Files Listing](#).
9. **Use the troubleshooting tools** described in the section below.
10. **Check for any improvement until the complete restoration of the system** after every step in the troubleshooting process.

If the problem remains, do not hesitate to contact the support team with all the information you have collected. The set of files needed include: the network capture file, the troubleshooting dump file and the last system backup.

Troubleshooting Tools

SOLIDserver provides users of the group *admin* with two ways of analyzing the system in case of a crash:

- The [Network Capture](#), that indicates the DHCP or DNS traffic on a given duration.
- The [Troubleshooting Dump](#), that allows to retrieve key debug information.

Both methods are complementary.

Network Capture

The network capture tool allows to capture packets on a given duration, i.e. the actions made through the appliance interface(s), to analyze DHCP and DNS traffic. When you run this utility, the archive file containing all the traffic information should be available in the directory listing module in the .pcap format.

To perform a network capture

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Network capture**. The wizard **Perform a network capture** opens.
3. In the drop-down list **Predefined**, you can select one of three options described in the table below:

Table 92.3. Available options in the field Predefined

Option	Description
Custom	Select this option to scan both the DNS and DHCP traffic. <i>Custom</i> is selected by default.
DHCP traffic	Select this option to scan the DHCP traffic. The field <i>Port</i> is automatically filled with 67.
DNS traffic	Select this option to scan the DNS traffic. The field <i>Port</i> is automatically filled with 53.

4. In the drop-down list **Interface**, select the interface for which you want to capture packets. It can either be *DEFAULT_INTERFACE* or the *DHCP_INTERFACE*.
5. In the field **Port**, you can specify the port for which you want to capture packets.
6. In the field **IP address**, you can specify the IP address for which you want to capture packets.
7. In the drop-down list **Protocol**, you can specify the protocol, either *udp*, *tcp* or both (*Any*).
8. In the drop-down list **Duration**, you can specify the duration of the capture, either *10s*, *30s*, *1mn*, *2mn* or *5mn*.
9. Click on **OK** to complete the operation. The report opens and closes. The page **Admin Home** is visible again.

The .pcap file containing all the traffic information is available on the page **Local files listing** accessible from the section **Maintenance** on the page **Admin Home**.

Troubleshooting Dump

The troubleshooting dump is a file containing DNS, DHCP and system debug data. Dumping troubleshooting details generates a *.tbz* archive file containing all the debug information. This file is available on the page *Local files listing*.

To generate and download the troubleshooting dump file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. **Generate the troubleshooting dump**
 - a. In the section **Expert**, click on **Troubleshooting dump**. The wizard **Troubleshooting dump** opens.
 - b. Tick the box(es) **Retrieve DNS information**, **Retrieve DHCP information** and/or **Retrieve system information**. The information ticked is included in the generated file.
 - c. Click on to complete the operation. The report opens and closes. The page **Admin Home** is visible again..
3. **Download the file**
 - a. In the section **Maintenance**, click on **Local files listing**. The page opens.
 - b. Filter the column **Name** with the keyword *DEBUG* to only list troubleshooting dump files.

If a *DEBUG* file does not have the extension *.tbz*, it means it is not fully generated yet.
 - c. Click on the name of the file of your choice to download it.

Managing Backups and Restoring Configurations

EfficientIP recommends that you regularly backup SOLIDserver. In order to help you perform this maintenance operation, SOLIDserver includes automatic backup and version management mechanism. The backup process can either be scheduled or triggered on demand.

The backup files are stored on the appliance itself, but you can also decide to store the backup files on a remote FTP server or SFTP server. For ease of use and to prevent confusion, binaries, system and log files are not included in the backup stored on the appliance. Still, they can be restored separately, either when you reinstall SOLIDserver or when you update the system.

DNS, DHCP and System logs can be included in the backup created on the remote archive.

Note that, when you save a backup of DNS zones managed via a smart architecture, you can later chose to restore them while keeping the latest version of the records they contain. You can also discard the latest changes and restore the records as saved in the backup file.

SOLIDserver automatically generates a new backup before each upgrade to allow reverting back its data and configuration.

Browsing the Backup Database

To display the list of backup files

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.

The page contains two panels:

Local backup files

This panel contains:

- The list of the backup files available on the appliance.
- The field **Hour of backup (SOLIDserver system time)**: the time of the daily backup.
- The field **Retention duration**: the number of days beyond which a backup is automatically deleted from the local database.

Remote archive

If the remote archiving is enabled, this panel contains the fields:

- **Remote server**: the address or hostname of the remote server storing the backup files.
- **Remote port**: the port number used on the remote server that communicates with SOLIDserver.
- **Remote directory**: the directory on the remote server where the backup files are stored.
- **Remote login**: the login used to connect to the remote FTP server.
- **Mode**: the protocol and mode used to connect to the remote server, either *Active FTP*, *Passive FTP* or *SFTP*.
- **Log DNS**: indicates if the DNS logs are included in the remote backup (*yes*) or not (*no*).
- **Log DHCP**: indicates if the DHCP logs are included in the remote backup (*yes*) or not (*no*).
- **System Log**: indicates if the System logs are included in the remote backup (*yes*) or not (*no*).
- **Retention duration**: the number of days beyond which a backup is automatically deleted from the remote server.

Creating an Instant Backup

You can create an instant backup of the whole system configuration on demand. An image of the system is generated and stored on the appliance. Each image can be then used to store the configuration of a SOLIDserver, which allows you to reload a previous backup in case of a revert back procedure.

Keep in mind that **creating an instant backup during the enrollment of a Hot Standby** appliance in High Availability **may trigger an error**.

To create an instant backup

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the menu, select **Tools > Create instant backup**. The wizard **Create instant backup** opens.
4. Tick the box **Exclude all the reports** if you want to exclude all generated HTML and PDF reports and only save the database, configuration and certification files.
5. Click on **OK** to complete the operation. The report opens and works for a while. Once the backup is generated, it is listed in the panel **Local backup file** and named *solid-<hostname>-<year><month><day>-<hour><minutes>.gz*.

Once generated, you can download your backup if need be. For more details, refer to the section [Downloading a Backup File](#).

Editing the Backup Settings

By default, SOLIDserver is configured to generate a backup of its database and network configuration every night at 2 a.m. This backup process can be edited to match your own schedule: it can run at a different hour, store the backup file for a limited time or, on the contrary, an unlimited number of days...

Editing the backup settings can maximize the disk space of your appliance if you schedule a backup rotation: the automatic deletion of obsolete backup files.

To schedule a daily backup

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the menu, select **Edit > Local backup files**. The wizard **Archive backup parameters** opens.
4. In the drop-down list **Hour of backup (SOLIDserver system time)**, select at what time you want to generate the daily backup. By default, it set to *1:00*.
5. You can tick the box **Exclude all the notifications** if you do not want to save the operation notifications in the daily backup file. These notifications are all listed in the window *Notifications* located left of the field *Global search*.
6. Click on **OK** to complete the operation.

To set a backup rotation

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the menu, select **Edit > Local backup files** or in the panel **Local backup files**, click on **EDIT**. The wizard **Archive backup parameters** opens.
4. In the drop-down list **Retention**, select the number of days beyond which a backup should be automatically deleted.
5. Click on **OK** to complete the operation.

Archiving the Backup Files on an FTP or SFTP server

You can archive a copy of SOLIDserver backup files on a remote server, FTP or SFTP.

During the configuration of the remote server, you can decide to include the DNS, DHCP and System logs or even specify a number of days beyond which they should be automatically deleted from the server. Note that:

- You can archive backups on an FTP server - via *Active FTP* or *Passive FTP* - or an SFTP server. We strongly recommend using SFTP which is far more secure than FTP as it uses an SSH key instead of a password.
- You can specify the port configured on the remote server. On SFTP servers, usually the same port than SSH is used.
- If no remote archive is configured, the panel **Remote archive** contains the message *Remote archive is disabled*.

To configure the remote FTP archive

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the panel **Remote archive**, click on . The wizard **Archive server parameters** opens.
4. Tick the box **Enable remote archive**, the wizard refreshes and displays the remote archive configuration parameters. By default, the box is unticked.
5. Configure the remote archive parameters according to your needs:

Table 92.4. Backup archiving parameters

Field	Description
Remote server	The IP address or the host name of the FTP or SFTP server.
Remote port	The port used to communicate with the server. If no port is used, the port 21 is used for FTP and the port 22 for SFTP.
Remote directory	The directory where the backup files should be stored.
Mode	The protocol and mode used to archive the files: <i>Active FTP</i> , <i>Passive FTP</i> or <i>SFTP</i> .
Remote login	The login of the account used to connect to the FTP or SFTP server.
Remote password	The password of the account used to connect to the FTP server. No password is required for SFTP.
DNS	Tick the box DNS, DHCP and/or System if you want to save the corresponding logs on the remote server.
DHCP	
System	
Retention	Select the number of days, from <i>4 days</i> to <i>Unlimited</i> , beyond which a backup should be automatically deleted from the FTP server. By default, <i>4 days</i> is selected.

6. Click on to complete the operation. The report opens and closes. The page refreshes and the panel **Remote archive** displays the FTP or SFTP server parameters you just configured.
7. If you selected SFTP, the panel **SSH local key** displays the SSH public key used. You must it and paste in on the SFTP server to secure the communication with SOLIDserver.

To disable the remote FTP archive

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the panel **Remote archive**, click on **[EDIT]**. The wizard **Archive server parameters** opens.
4. Untick the box **Enable remote archive**, the wizard refreshes. It is now empty.
5. Click on **[OK]** to complete the operation. The report opens and closes. The page refreshes and the panel **Remote archive** contains the message *Remote archive is disabled*.

Downloading a Backup File

At any time administrators can download a backup file.

To download a backup file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the panel **Local backup file**, select the backup file of your choice.
4. You can click on **[DOWNLOAD]** to save the file locally, this automatically closes the wizard. The page **Backup & Restore** is visible again.

Uploading a Backup File

At any time administrators can upload a backup file located outside SOLIDserver to the local SOLIDserver file system from the GUI.

To upload a backup file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the menu, select ✨. **Tools > Upload a backup file**. The wizard **Upload SOLIDserver backup** opens.
4. Click on **[BROWSE]** to select the image to upload from your local file system.
5. Click on **[OK]** to complete the operation. The backup file is now listed in the panel **Local backup files**.

Restoring a Backup File

You can restore SOLIDserver database and configuration from a backup file. Keep in mind that:

- **You need the backup file name and version number.** That's why each new backup generates an increment number that concatenates the date and hour as follows: *solid-<hostname>-<year><month><day>-<hour><minutes>.gz*.

If the backup file you need is not listed, upload it as detailed in the section [Uploading a Backup File](#).

- **The backup file contains the appliance data and configuration** as they were set at the time of the backup generation. During the restoration, you can choose to include both or only the data:
 - The **appliance data** includes objects from all modules including the ones of the *Local Files Listing* as well as all the rules, notifications and reports, unless you excluded them when you saved the backup.
 - The **appliance configuration** includes the network configuration (hostname, DNS resolver, firewall configuration, default gateways, default/static route configuration) and services configuration (services status, xfer account settings, SNMP communities).
- **The records belonging to a smart architecture have a dedicated restoration option.** If your backup file includes zones belonging to a physical server managed via a smart architecture, you can decide to:
 - Overwrite their content with the one saved in the backup file. In this case, you lose all the latest changes (records added or deleted) between the time of the backup and the time of the restoration.
 - Keep their current content if they are present in the backup file.

In both cases, all the zones that are not part of the backup file, as well as their records, are lost during the restoration.

- **You cannot restore a backup on an appliance set in High Availability.** You need to disable the High Availability, restore the backup on a Standalone appliance and then configure the High Availability again. For more details, refer to the section [Replacing a Hot Standby Appliance With Backup](#).
- **Restoring objects of the module Application overwrites the current database.** If you created, edited or deleted objects since the backup was saved, all changes are lost and may even not be visible in the GUI. Before restoring a backup, make sure you saved it when the Application database was up-to-date.
- **Restoring objects of the module Guardian overwrites information displayed in the GUI.** The information displayed on the pages *All policies* and *All triggers* might be erroneous but the configuration on the server is still correct and takes into account everything that you created, edited and/or deleted since the backup was saved.

After the backup restoration, everything that you change in the GUI is pushed on the server and overwrites its current configuration.

To restore a backup file

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
3. In the panel **Local backup files**, select the backup file you want to restore.
4. Click on [RESTORE](#). The wizard **Restore a backup file** opens.
5. Tick the box **Restore the system configuration** to restore the backup system configuration. Tick it if you are restoring a backup using an NSD or Unbound Hybrid server. Otherwise, the backup data is restored but the current system configuration of the appliance is kept.

6. Tick the box **Overwrite DNS records managed via a smart architecture** to restore the records database as saved during the backup. If you do not tick the box, the restored zones keep the current version of the records they contain if they are managed via a smart architecture.

In both cases, a restoration includes all and only the zones present in the backup file.

7. Click on to complete the operation.

Shutting Down and Rebooting

When you reboot or shut down an appliance, SOLIDserver transfers file system cache to disk, stops all running processes and then reboots or halts (shutdown).

Note that once SOLIDserver operating system is stopped, the power supplies are automatically turned off.

Rebooting SOLIDserver

There are two ways of rebooting SOLIDserver safely, from the GUI and via CLI.

Note that **only users with sufficient rights can reboot SOLIDserver**.

Rebooting SOLIDserver From the GUI

Only users with sufficient rights can reboot any appliance from the GUI.

To reboot the system from the GUI

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Reboot the system**. The wizard **Reboot the system** opens.
3. Click on to complete the operation.

Rebooting SOLIDserver Through CLI

Only users with sufficient rights can reboot any appliance through CLI using its IP address, hostname or a serial port.

To reboot the system from the CLI

1. Connect to SOLIDserver CLI via a shell session or a port console.
2. Once you are connected, the page **WELCOME TO SOLIDSERVER** opens.
3. Log in using the credentials *admin/admin*. The **Main menu** appears.
4. Hit the key **P** key to select **P Power Management**.
5. Hit the key **Enter**, the page **Power Management** opens.
6. Hit **R** to select **R Reboot SOLIDserver**.

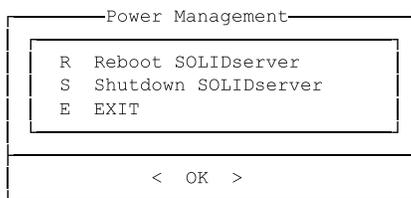
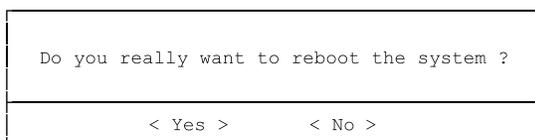


Figure 92.1. Power management

Hit **Enter**, the confirmation window opens.

7. The button **Yes** is highlighted.



```
*** FINAL System shutdown message from admin@solid.intranet ***
System going down IMMEDIATELY
```

Figure 92.2. Reboot confirmation and final message

Hit **Enter** to reboot SOLIDserver software.

Shutting Down SOLIDserver

SOLIDserver is designed to operate continuously, so under normal circumstances, you do not need to turn it off or shut it down. However, if you have to turn it off, you can use the GUI, CLI or the hardware appliance itself.

Prerequisites

- Only users with sufficient rights can shut down SOLIDserver software.
- Before shutting down a remote SOLIDserver hardware appliance, keep in mind that you must have physical access to it power it back with the the button . You must plug in again SDS-50 appliances to start them back.

Shutting Down SOLIDserver From the GUI

Any SOLIDserver appliance can be shut down from the GUI, granted that the user has sufficient rights.

To shut down SOLIDserver from the GUI

1. Make sure you meet the [Prerequisites](#).
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Maintenance**, click on **Shutdown the system**. The wizard **Shutdown the system** opens.
4. Click on  to complete the operation.

Shutting Down SOLIDserver Via CLI

Any SOLIDserver appliance can be shut down via CLI using its IP address, hostname or a serial port.

To shut down SOLIDserver via CLI

1. Make sure you meet the [Prerequisites](#).
2. Connect to SOLIDserver CLI via a shell session or a port console.
3. Log in using the credentials *admin/admin*. The **Main menu** appears.
4. Hit the key **P** to select **P Power Management**.

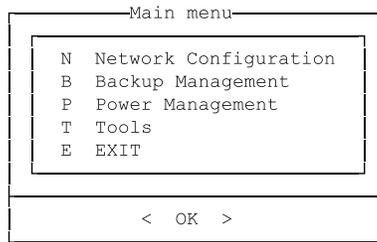


Figure 92.3. Main menu

Hit the key **Enter**, the page **Power Management** opens.

5. Hit **S** to select **S Shutdown SOLIDserver**.

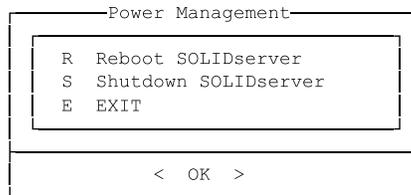


Figure 92.4. Power management

Hit **Enter**, the confirmation window opens.

6. The button **Yes** is highlighted.

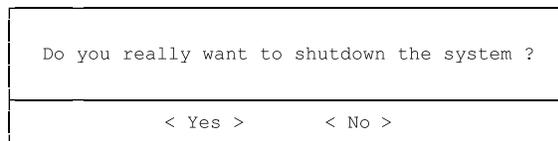


Figure 92.5. Shutdown confirmation

Hit **Enter** to shutdown SOLIDserver software.

Shutting Down SOLIDserver From the Hardware Itself

Any SOLIDserver hardware appliance can be shut down for the hardware itself, to avoid an abrupt shutdown follow the procedure below.

It must be done as a last resort if it is impossible from the GUI or via CLI, note that shutting down the hardware appliance also shuts down the software.

To shut down SOLIDserver from the hardware itself

1. Make sure you meet the [Prerequisites](#).
2. From the front panel of the appliance, quickly press and release the power button .

Note that a long press stops the appliance abruptly and may trigger errors. The appliance would stop without properly transferring files or stopping the processes as expected.

For SDS-50 appliances, you must unplug the appliance and plug it again.

3. The appliance stops automatically after synchronizing its buffer on the disk.

Chapter 93. Upgrading

This chapter details the procedures to successfully upgrade SOLIDserver 7.1 to a higher patch.

To upgrade SOLIDserver from a previous minor or major version, refer to the guide *SOLIDserver_Upgrade_to_Version_7.1.pdf* available on our website¹.

The upgrade process can take a while as the appliance:

1. Generates and saves a backup of the database at the time of the upgrade before rebooting.
2. Upgrades the appliance version and database schema before rebooting a second time.

Prerequisites

1. **Have an Internet connection and your credentials ready** to download the version of SOLIDserver that suits your needs on our website.
2. **If you installed hot-fixes**, they are deleted during the upgrade. The upgrade wizard retrieves a list of all installed hot-fixes that you can download.

Note that hot-fixes are only detected locally. When you upgrade a remote SOLIDserver from the Management appliance, hot-fixes are deleted but the list of files is not available.

3. **Keep the upgrade file as is.** You cannot rename SOLIDserver upgrade files, otherwise the upgrade may fail.
4. **Follow the proper upgrade procedure.**
 - If you want to upgrade an appliance, either managing remote appliances or not, refer to the section [Upgrading an Appliance](#).
 - If you want to upgrade appliances managed remotely, refer to the section [Upgrading Appliances Managed Remotely](#).
 - If you want to upgrade appliances configured in High Availability, refer to the section [Upgrading Appliances in High Availability](#).
5. **Save the backup file generated during the upgrade.** All backup files are automatically deleted after the number of days defined in the *Retention duration* you set. If for any reason you need to troubleshoot the upgrade beyond that period, you will need the latest backup. For more details regarding troubleshooting, refer to the section [Troubleshooting the Upgrade](#) below.

For more details regarding the backup retention time, refer the section [Managing Backups and Restoring Configurations](#).

Upgrading an Appliance

If you meet the [prerequisites](#), you can upgrade SOLIDserver to a higher patch, it implies to:

1. Download the image of the latest patch of version 7.1.
2. Upgrade the appliance.
3. Save the backup generated during the upgrade.

¹At <https://downloads.efficientip.com/support/downloads/SOLIDserver/7.1/docs/>, log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

Note that:

- Once the appliance is upgraded, the page *Services configuration* might display different statuses: any service that was stopped restarts when SOLIDserver reboots. Only disabled services are not started after an upgrade.
- If an error occurs during the upgrade, refer to the section [Troubleshooting the Upgrade of a Standalone or Remote Appliance](#).
- If you are upgrading an appliance in High Availability or managed remotely, refer to the section [Upgrading Appliances in High Availability](#) or [Upgrading Appliances Managed Remotely](#).

To upgrade a Standalone or Management appliance

Only users of the group *admin* can perform this operation.

1. Download the image of the latest patch of version 7.1

- a. Use your client account credentials to connect to the download portal: <https://downloads.efficientip.com/support/downloads/SOLIDserver/>.
- b. Open the folder *7.1/*.
- c. Download the file *solidserver-<architecture>-<version>* that suits your needs. The file must be without extension and its version must be higher than the version you are upgrading from.

2. Upgrade the appliance

- a. Connect to the appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Upgrade**. The wizard **Upgrade SOLIDserver** opens.
- d. Click on **BROWSE** to select the file containing the upgrade image. The name of the selected file is displayed in the field **File name**.

Once the file retrieved, the fields **Current version** and **Upgrade to** appear. They both indicate the version and architecture as follows: *<version>(<architecture>)*.

If you installed hot-fixes, all related files are listed under *Modified files* on the page. For more details, refer to the [prerequisites](#).

You can click on **DOWNLOAD FULL REPORT** to retrieve both lists.

- e. Click on **NEXT**. The last page of the wizard opens.
- f. To start the upgrade, click on **UPGRADE**.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process**. Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

3. Save the backup file

- a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.

- b. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- c. In the panel **Local backup file**, select the latest backup file. It is named `solid-<hostname>-<year><month><day>-<hour><minutes>.gz`.
- d. Click on **DOWNLOAD** to save the file locally. Once the file is downloaded, the page is visible again.

Upgrading Appliances Managed Remotely

If you meet the [prerequisites](#), you can upgrade SOLIDserver appliances managed remotely to a higher patch.

You can upgrade all remote appliances at once to the same version and architecture as their Management appliance. Which is why you must:

1. Check the version of your Management appliance.
2. Upgrade the remote appliance(s) to the same version from the Management appliance.
3. Save the backup file of each remote appliance, it is generated during the upgrade.

Note that:

- **To upgrade a remote appliance, the Management appliance should already be upgraded.**
- When you upgrade remote appliances from the Management appliance, you cannot retrieve the list of hot-fixes because it is only available locally.
- Once appliances are upgraded, the page *Services configuration* might display different statuses: any service that was stopped is started when SOLIDserver reboots. Only disabled services are not started after an upgrade.
- If an error occurs during the upgrade, refer to the section [Troubleshooting the Upgrade of a Standalone or Remote Appliance](#).
- If you are upgrading a Standalone appliance or only the Management appliance, refer to the section [Upgrading an Appliance](#).
- If you are upgrading appliances in HA, refer to the section [Upgrading Appliances in High Availability](#).

To upgrade one or several remote appliances

Only users of the group *admin* can perform this operation.

1. **Check the version of the Management appliance**
 - a. Connect to the **Management appliance** GUI.
 - b. On the **Main dashboard**, the gadget **System Information** indicates the appliance version and architecture. Once upgraded, the remote appliance(s) should match this version.
2. **Upgrade the remote appliance(s) from the Management appliance**
 - a. Connect to the **Management appliance** GUI.
 - b. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.

- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. Tick the appliance(s) you want to upgrade.
- e. In the menu, select  **Edit** > **Upgrade remote SOLIDserver**. The wizard **Upgrade remote SOLIDserver** opens.

If you installed hot-fixes, you cannot retrieve the list of these files as they are only available locally.

- f. To start the upgrade, click on **OK**. The report opens and works until the selected appliance(s) version matches the version of the Management appliance. The wizard eventually closes. The appliance(s) are not accessible for a few minutes.

3. Save the backup file of each remote appliance

- a. Connect to the **remote appliance** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- d. In the panel **Local backup file**, select the latest backup file. It is named *solid-<hostname>-<year><month><day>-<hour><minutes>.gz*.
- e. Click on **DOWNLOAD** to save the file locally. Once the file is downloaded, the page is visible again.
- f. Repeat steps a to e for each remote appliance you upgraded.

Upgrading Appliances in High Availability

There are two ways of upgrading appliances configured in High Availability, together or separately.

In addition to the basic upgrade [prerequisites](#), you must go to the page *Centralized Management* of both appliances to **make sure that the High Availability configuration is properly set**:

- One appliance must be the **Master** and have the Status  **OK**.
- The other appliance must be the **Hot Standby** and have the Status  **Managed (remote)**.
- Both appliances must share the **same HA UID**.
- Both appliances must share the **same version**.
- There should be **little to no Time drift** between them.
- The page *Centralized Management* of the Master and the Hot Standby appliances must display the same information.

Note that:

- Right before the upgrade itself, we disable the automatic re-enrollment option to avoid errors. This prevents the Master appliance from re-enrolling the Hot Standby if the upgrade takes too long. After the upgrade, you can enable it again.

- Once the appliances are upgraded, the page *Services configuration* might display different statuses: any service that was stopped is started when SOLIDserver reboots. Only disabled services are not started after an upgrade.
- If an error occurs during the upgrade, refer to the section [Troubleshooting the Upgrade of Appliances in High Availability](#).
- If you are not upgrading an appliance in High Availability, refer to the section [Upgrading an Appliance](#) or [Upgrading Appliances Managed Remotely](#).

You can upgrade appliances in High Availability in two different ways:

1. **Upgrade both appliances at once from the Master appliance.** For more details, refer to the section [Upgrading Both Appliances at Once](#).
2. **Upgrade the appliances one after the other, starting with the Hot Standby appliance.** For more details, refer to the section [Upgrading One Appliance at a Time](#).

Upgrading Both Appliances at Once

You can upgrade both appliances in High Availability at once from the Master appliance. This process ensures that:

1. **The Hot Standby appliance is upgraded first.** As the upgrade requires to stop and restart an appliance that would imply switching the appliances role, if the Hot Standby is upgraded first, the Master appliance database is still available and no switch is required.
2. **The Master appliance is upgraded once the Hot Standby upgrade is complete.** Once the Hot Standby is upgraded, the Master appliance can be stopped and restarted and no switch is performed.

In High Availability, upgrading both appliances at once must follow this order because upgrading an appliance stops and restarts it. This process ensures that the appliances do not switch roles and that the database is available even during the upgrade.

Therefore, from the Master appliance you can safely upgrade both appliances as detailed in the procedure below.

To upgrade both HA appliances at once

Only users of the group *admin* can perform this operation.

1. **Make sure both appliances meet the prerequisites**

The basic upgrade [prerequisites](#) and the [High Availability prerequisites](#).

2. **Download the image of the latest patch of version 7.1**

- a. Use your client account credentials to connect to the download portal: <https://downloads.efficientip.com/support/downloads/SOLIDserver/>.
- b. Open the folder *7.1/*.
- c. Download the file *solidserver-<architecture>-<version>* that suits your needs. The file must be without extension and its version must be higher than the version you are upgrading from.

3. Update the HA files database from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. In the menu, select **Tools > Update HA files database**. The report opens.
- e. Click on **OK** to complete the operation. The report opens and closes. The page is visible again.

4. Disable the automatic re-enrollment from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- d. Filter the column **Name** with the keyword `module.system.auto_replication_repair` to display the entry.
- e. In the column **Value**, if the entry is set to `0`, the automatic re-enrollment is already disabled. Go straight to step 5 to upgrade your appliances.
- f. In the column **Value**, if the entry is different from `0`, click on said value. The related wizard opens.
- g. In the field **Value**, specify `0`.
- h. Click on **OK** to complete the operation. The wizard reloads and closes.

5. Upgrade to version 7.1 both appliances from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Upgrade**. The wizard **Upgrade SOLIDserver** opens.
- d. Click on **BROWSE** to select the file containing the upgrade image. The name of the selected file is displayed in the field **File name**.

Once the file retrieved, the fields **Current version** and **Upgrade to** appear. They both indicate the version and architecture as follows: `<version>(<architecture>)`.

If you installed hot-fixes on the Master appliance, all related files are listed under *Modified files* on the page. For more details, refer to the [prerequisites](#).

You can click on **DOWNLOAD FULL REPORT** to retrieve both lists.

- e. Click on **NEXT**. The last page of the wizard opens.
- f. To start the upgrade, click on **UPGRADE**.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process.** Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

6. Check that both appliances are upgraded

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. In the column **Version**, the value for the **Master** should match the one on the **Hot Standby** appliance.
- e. Connect to the **Hot Standby** appliance GUI and repeat steps *b* to *d* to make sure both appliances are in the same version of SOLIDserver. If they do not, refer to the section [Troubleshooting the Upgrade of Appliances in High Availability](#).

7. Save the backup file of both appliances

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- d. In the panel **Local backup file**, select the latest backup file. It is named *solid-<hostname>-<year><month><day>-<hour><minutes>.gz*.
- e. Click on **DOWNLOAD** to save the file locally. Once the file is downloaded, the page is visible again.
- f. Connect to the **Hot Standby** appliance GUI and repeat the steps *b* to *e* if you set a specific network/services configuration.

8. Enable the automatic re-enrollment from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on ⚙ **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- d. Filter the column **Name** with the keyword *module.system.auto_replication_repair* to display the entry.
- e. In the column **Value**, click on *0*. The related wizard **Registry database Edit a value** opens.
- f. In the field **Value**, specify *1*.
- g. Click on **OK** to complete the operation. The wizard reloads and closes.

Upgrading One Appliance at a Time

Even if the upgrade is automated by default, you can still upgrade one appliance after the other. To ensure that the database is available throughout the upgrades, you must:

1. **Upgrade the Hot Standby appliance.** The Hot Standby must be upgraded first to make sure the Master database is available.
2. **Upgrade the Master appliance.** Once the Hot Standby appliance upgrade is complete, you can upgrade the Master. This avoids a switch and ensures the database availability.

To upgrade one HA appliance at a time

Only users of the group *admin* can perform this operation.

1. Make sure both appliances meet the prerequisites

The basic upgrade [prerequisites](#) and the [High Availability prerequisites](#).

2. Download the image of the latest patch of version 7.1

- a. Use your client account credentials to connect to the download portal: <https://downloads.efficientip.com/support/downloads/SOLIDserver/>.
- b. Open the folder *7.1/*.
- c. Download the file *solidserver-<architecture>-<version>* that suits your needs. The file must be without extension and its version must be higher than the version you are upgrading from.

3. Update the HA files database from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. In the menu, select  **Tools** > **Update HA files database**. The report opens.
- e. Click on to complete the operation. The report opens and closes. The page is visible again.

4. Disable the automatic re-enrollment from the Master

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
- d. Filter the column **Name** with the keyword *module.system.auto_replication_repair* to display the entry.
- e. In the column **Value**, if the entry is set to *0*, the automatic re-enrollment is already disabled. Go straight to step 5 to upgrade your appliances.

- f. In the column **Value**, if the entry is different from *0*, click on said value. The related wizard opens.
- g. In the field **Value**, specify *0*.
- h. Click on **OK** to complete the operation. The wizard reloads and closes.

5. **Upgrade to version 7.1 the appliances one after the other**

- a. Connect to the **Hot Standby** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Upgrade**. The wizard **Upgrade SOLIDserver** opens.
- d. Click on **BROWSE** to select the file containing the upgrade image. The name of the selected file is displayed in the field **File name**.

Once the file retrieved, the fields **Current version** and **Upgrade to** appear. They both indicate the version and architecture as follows: *<version>(<architecture>)*.

If you installed hot-fixes, all related files are listed under *Modified files* on the page. For more details, refer to the [prerequisites](#).

You can click on **DOWNLOAD FULL REPORT** to retrieve both lists.

- e. Click on **NEXT**. The last page of the wizard opens.
- f. To start the upgrade, click on **UPGRADE**.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process**. Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

- g. Connect to the **Master** appliance GUI and repeat steps *b* to *f* to upgrade it. During the upgrade, the Master appliance switches role with the Hot Standby.

6. **Check that both appliances are upgraded**

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. In the column **Version**, the value for the **Master** should match the one on the **Hot Standby** appliance.
- e. Connect to the **Hot Standby** appliance GUI and repeat steps *b* to *d* to make sure both appliances are in the same version of SOLIDserver. If they do not, refer to the section [Troubleshooting the Upgrade of Appliances in High Availability](#).

7. **Save the backup file of both appliances**

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

- c. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
 - d. In the panel **Local backup file**, select the latest backup file. It is named *solid-<hostname>-<year><month><day>-<hour><minutes>.gz*.
 - e. Click on **DOWNLOAD** to save the file locally. Once the file is downloaded, the page is visible again.
 - f. Connect to the **Hot Standby** appliance GUI and repeat the steps *b* to *e* if you set a specific network/services configuration.
8. **Enable the automatic re-enrollment from the Master**
- a. Connect to the **Master** appliance GUI.
 - b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
 - c. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
 - d. Filter the column **Name** with the keyword *module.system.auto_replication_repair* to display the entry.
 - e. In the column **Value**, click on *0*. The related wizard **Registry database Edit a value** opens.
 - f. In the field **Value**, specify *1*.
 - g. Click on **OK** to complete the operation. The wizard reloads and closes.

Troubleshooting the Upgrade

To troubleshoot an upgrade you can roll back to the version previously installed:

1. Connect to the appliance via CLI.
2. Execute the command that rolls back the installation and restores the backup saved right before the upgrade.
3. Wait for the appliance to reboot.

Follow the proper troubleshooting procedure:

- The procedure is the same for Standalone, Management or remote appliances. For more details, refer to the section [Troubleshooting the Upgrade of a Standalone or Remote Appliance](#).
- To troubleshoot an appliance configured in High Availability refer to the section [Troubleshooting the Upgrade of Appliances in High Availability](#).

Troubleshooting the Upgrade of a Standalone, Management or Remote Appliance

If the appliance was upgraded to the correct version but an error occurred, whether it is remotely managed or not, you can roll back via CLI. This restores the appliance version and database as they were before the upgrade.

To troubleshoot an appliance configured in High Availability refer to the section [Troubleshooting the Upgrade of Appliances in High Availability](#).

Note that you need the backup file of the previous version or patch to be stored locally on the appliance. It was generated during the upgrade to version 7.1.

To roll back the upgrade of a Standalone, Management or remote appliance

Only users of the group *admin* can perform this operation.

1. Make sure the backup file is stored on the appliance

- a. Connect to the appliance GUI.
- b. On the **Main Dashboard**, in the gadget **System information**, check the **Version** currently installed on the appliance.
- c. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- d. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- e. In the panel **Local backup files**, make sure the backup generated during the upgrade to the desired version is listed.

If the backup file is not listed, in the menu select **Tools > Upload a backup file**. Browse your computer to find it and click on **OK** to make sure it is listed in the panel.

2. Roll back the appliance

- a. Using a terminal emulator, open a shell session to connect to the appliance CLI you are troubleshooting using its IP address or hostname.
- b. Log in using the credentials *admin/admin*. The CLI **Main menu** appears.

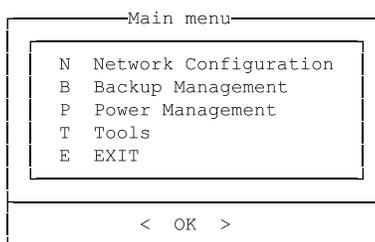


Figure 93.1. Main menu

- c. Hit the key **T** to select **T Tools** and hit the key **Enter**. The page **Tools** opens.
- d. The line **S Start a shell** is selected, hit **Enter**. The terminal closes.
- e. Execute the following command to use *root* permissions:

```
% su
```

Hit **Enter**. The % changes to a #.

- f. Execute the rollback command:

```
# /usr/local/nessy2/script/rollback_upgrade.sh
```

The menu **SOLIDserver rollback** opens.

- g. In the menu **Choose backup to restore to <previous-version>-<architecture>**, use the arrows of your keyboard to highlight the backup file that suits your needs.

Note that only the backup files saved when the previous version was installed are listed, so if you upgraded a few days ago, none of the backup files saved between the upgrade and the rollback are listed.

- h. Once the backup file of your choice is highlighted, click on **Enter**. The page **WARNING ROLLBACK WILL RESTORE A BACKUP** opens.
- i. Hit the key **Y** to highlight the line **() Y CONFIRM ROLLBACK**. Press the key **Space** to select this option, the ***** indicates the line is selected: **(*) Y CONFIRM ROLLBACK**.
- j. Hit **Enter** to confirm.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process**. Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

Troubleshooting the Upgrade of Appliances in High Availability

To troubleshoot the upgrade of appliances in High Availability, check if both appliances upgraded to the same version and then follow one of these two procedures:

- [If HA Appliances Were Both Upgraded to the Same Version.](#)
- [If HA Appliances Were Both Upgraded to Different Versions.](#)

If both appliances were correctly upgraded to the desired version but the High Availability configuration was broken, you only need to reconfigure it. For more details, refer to the section [Setting a High Availability Configuration](#).

If HA Appliances Were Both Upgraded to the Same Version

If an error occurred but both HA appliances were upgraded to the same patch of version 7.1, you can roll back via CLI starting with the Master appliance. This restores the appliances version and database as they were before the upgrade.

Before troubleshooting HA appliances:

- You need the backup file, generated during the upgrade, to be stored locally on each appliance.
- You need to disable the high availability configuration, if it is enabled, before rolling back.

To roll back the upgrade of HA appliances

Only users of the group *admin* can perform this operation.

1. **Make sure the backup file is stored on both appliances**
 - a. Connect to the **Master** appliance GUI.
 - b. On the **Main Dashboard**, in the gadget **System information**, check the **Version** currently installed on the appliance.
 - c. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

- d. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- e. In the panel **Local backup files**, make sure the backup generated during the upgrade to the desired version is listed.

If the backup file is not listed, in the menu select **Tools > Upload a backup file**. Browse your computer to find it and click on **OK** to make sure it is listed in the panel.

- f. Connect to the **Hot Standby** appliance GUI and repeat the steps *b* to *e*.

2. Make sure the High Availability is properly disabled

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.

If the High Availability configuration is already disabled, i.e. both appliances have a *Standalone* role, go to step 3.

- d. Tick the Hot Standby appliance.
- e. In the menu, click on **Delete**. The wizard **Delete** opens.
- f. Click on **OK** to complete the operation. The report opens and works for a while - the Hot Standby database is saved before it is erased - and closes. The Hot Standby is no longer listed on the page. The former Master appliance keeps a *Master* role.

The former Hot Standby appliance is not accessible while it is reset.

3. Roll back the appliances

- a. Using a terminal emulator, open a shell session to connect to the **Master** appliance CLI you are troubleshooting using its IP address or hostname.
- b. Log in using the credentials *admin/admin*. The CLI **Main menu** appears.

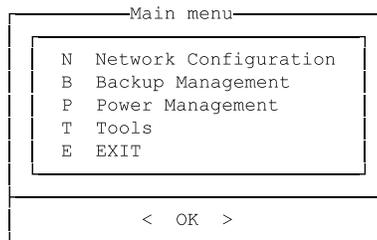


Figure 93.2. Main menu

- c. Hit the key **T** to select **T Tools** and hit the key **Enter**. The page **Tools** opens.
- d. The line **S Start a shell** is selected, hit **Enter**. The terminal closes.
- e. Execute the following command to use *root* permissions:

```
% su
```

Hit **Enter**. The % changes to a #.

- f. Execute the rollback command:

```
# /usr/local/nessy2/script/rollback_upgrade.sh
```

The menu **SOLIDserver rollback** opens.

- g. In the menu **Choose backup to restore to <previous-version>-<architecture>**, use the arrows of your keyboard to highlight the backup file that suits your needs.

Note that only the backup files saved when the previous version was installed are listed, so if you upgraded a few days ago, none of the backup files saved between the upgrade and the rollback are listed.

- h. Once the backup file of your choice is highlighted, click on **Enter**. The page **WARNING ROLLBACK WILL RESTORE A BACKUP** opens.
- i. Hit the key **Y** to highlight the line () **Y CONFIRM ROLLBACK**. Press the key **Space** to select this option, the * indicates the line is selected: **(*) Y CONFIRM ROLLBACK**.
- j. Hit **Enter** to confirm.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process**. Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

- k. Connect to the **Hot Standby** appliance CLI and repeat the steps *b* to *j*.

After the rollback, you might want or need to re-enroll the appliances to make sure there is no replication delay. For more details, refer to the section [Configuring High Availability Advanced Options](#).

If HA Appliances Were Both Upgraded to Different Versions

If both appliances in High Availability were upgraded to two different versions, a message under the menu prompts you to downgrade your software version. You can downgrade each appliance locally via CLI, starting with the Master appliance.

Note that:

- You need the backup file, generated during the upgrade, to be stored locally on each appliance.
- You need to disable the high availability configuration, if it is enabled, before downgrading both appliances separately.

To downgrade HA appliances upgraded to different versions

Only users of the group *admin* can perform this operation.

1. Download SOLIDserver image

- a. Use your client account credentials to connect to the download portal: <https://downloads.efficientip.com/support/downloads/SOLIDserver/>.
- b. Download the file *solidserver-<architecture>-<version>* of the last version installed on the appliances before the upgrade. The file must be without extension.

2. Make sure the backup file is stored on both appliances

- a. Connect to the **Master** appliance GUI.

- b. On the **Main Dashboard**, in the gadget **System information**, check the **Version** currently installed on the appliance.
- c. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- d. In the section **Maintenance**, click on **Backup & Restore**. The page **Backup & Restore** opens.
- e. In the panel **Local backup files**, make sure the backup generated during the upgrade to the desired version is listed.

If the backup file is not listed, in the menu select . *Tools > Upload a backup file*. Browse your computer to find it and click on *OK* to make sure it is listed in the panel.

- f. Connect to the **Hot Standby** appliance GUI and repeat the steps *b* to *e*.

3. Make sure the High Availability is properly disabled

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.

If the High Availability configuration is already disabled, i.e. both appliances have a *Standalone* role, go to step 3.

- d. Tick the Hot Standby appliance.
- e. In the menu, click on  **Delete**. The wizard **Delete** opens.
- f. Click on to complete the operation. The report opens and works for a while - the Hot Standby database is saved before it is erased - and closes. The Hot Standby is no longer listed on the page. The former Master appliance keeps a *Master* role.

The former Hot Standby appliance is not accessible while it is reset.

4. Downgrade both appliances separately

- a. Connect the future **Master** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **Maintenance**, click on **Upgrade**. The wizard **Upgrade SOLIDserver** opens.
- d. Click on to select the file containing the upgrade image. The name of the selected file is displayed in the field **File name**.

Once the file retrieved, the fields **Current version** and **Upgrade to** appear. They both indicate the version and architecture as follows: *<version>(<architecture>)*.

If you installed hot-fixes, all related files are listed under *Modified files* on the page. For more details, refer to the [prerequisites](#).

You can click on to retrieve both lists.

- e. Click on . The last page of the wizard opens.

- f. To start the upgrade, click on **UPGRADE**.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process.** Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

- g. Connect to the future **Hot Standby** appliance GUI and repeat the steps *b* to *f*.

5. Restore the backup of both appliances separately from the CLI

- a. Using a terminal emulator, open a shell session to connect to the **Master** appliance CLI using its IP address or hostname.

If the connection cannot be established, it means that SOLIDserver did not finish downgrading. Wait and open another shell session.

- b. Log in using the credentials *admin/admin*. The CLI **Main menu** appears.

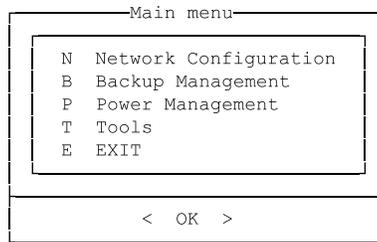


Figure 93.3. Main menu

- c. Hit the key **B** to select **B Backup Management** and hit the key **Enter**. The page **Backup Management** opens.
- d. Select **R Restore backup** and hit **Enter**. The page **Backup files** opens.
- e. Using the digit keys, select the backup file generated before the unsuccessful upgrade.
- f. Hit **Enter** to select it. The message **Do you also want to restore the configuration of the system?** appears.
- g. The answer **Yes** is selected. Hit **Enter** to select it.

The message **Do you really want to restore <backup-file-name>? Your server will automatically be rebooted to complete the process if you continue** appears and the answer **Yes** is selected

- h. Hit **Enter** to confirm.

After a while, the connection is automatically interrupted as the appliance shuts down and reboots. **Do not interrupt the process.** Once the GUI is reachable again, it displays the restart progression. When the login fields appear, you can connect to the appliance.

- i. Connect to the **Hot Standby** appliance CLI and repeat steps *b* to *h* if you want to restore the network/services configuration of the future Hot Standby appliance.

Note that, if you do not need to restore the network/services configuration of your Hot Standby appliance, you can restore only the Master appliance.

6. Enroll the Hot Standby again

- a. Connect to the **Master** appliance GUI.
- b. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- c. In the section **System**, click on **Centralized Management**. The page **Centralized Management** opens.
- d. In the menu, click on **+ Add**. The **Add/Modify remote SOLIDserver** appears.
- e. In the field **SOLIDserver IP address**, specify the IPv4 address of the appliance you want to add to the list.
- f. If the field **"Admin" account password** is empty, type in the SSH password, i.e. the default one (*admin*) or the one you set if you changed it.
- g. Click on  to complete the operation. The report opens and closes. The new appliance is listed and marked **Standalone** in the column *Role* and **Managed (remote)** in the column *Status*.
- h. Tick the future Hot Standby.
- i. In the menu, select  **Edit > Enroll SOLIDserver as Hot Standby**. The wizard **Enroll SOLIDserver as Hot Standby** opens.
- j. Click on  to complete the operation. The report opens and works for a while, until the **Hot Standby** appliance database is erased and replaced by the **Master** appliance database. The appliance set as **Hot Standby** is unavailable for a while. Each appliance role is modified according to the configuration, they both get the same HA UID.

7. **Contact your reseller's support team** to correct your database and successfully upgrade SOLIDserver.

After the rollback, you might want or need to re-enroll the appliances to make sure there is no replication delay. For more details, refer to the section [Configuring High Availability Advanced Options](#).

Part XIX. Customization

There are many ways of customizing SOLIDserver to suit your needs from the GUI to the data you manage.

You can **customize the GUI** following the chapters:

- [Customizing the GUI](#) details how to specify the text of your choice in the home page Welcome banner, or include images on the Home page or the login page.
- [Managing Smart Folders](#) details how to set up a tree-like display of your data for the modules IPAM, DHCP, DNS, NetChange and Device Manager using Smart Folders in the Tree View. This completely virtual organization of the objects can help you simplify deep organizations.
- [Managing IPv6 Labels](#) details how to create and use IPv6 labels above the IPv6 addresses of your IPAM and DHCP entries and differentiate containers more easily.

You can also **customize your databases** even further following the chapters:

- [Configuring Classes](#) details how administrators can manage classes to configure custom properties for your resources and tailor your database.
 - [Custom DB](#) details how to create databases tailored to your needs. They can, for instance, be used when to ease the configuration of your classes.
 - [Managing Customization Packages](#) details how to install packages to import customized functionalities from an archive file.
-

Chapter 94. Customizing the GUI

You can customize SOLIDserver GUI with images on the pages *Home* and *Login*, edit the welcome banner message or even edit the fields' name to suit your needs.

Customizing SOLIDserver Login Page With an Image

SOLIDserver provides the possibility to display an image on the appliance Login page: you can either display it over the logo of the login window or as a background.

At any time, this image can be changed or removed. Only users of the group *admin* can perform these changes.

To customize the login page with an image you need to:

1. Upload the image to the local files listing;
2. Specify the image name as the value of the dedicated registry database item.

Uploading an Image to SOLIDserver

You can upload any image to SOLIDserver database. Keep in mind that uploading images with a transparent background allows to fully integrate them to SOLIDserver graphical interface.

The login page image maximum size is 373x74 pixels. You can also upload a thumbnail of 45 pixels in height to be displayed in the upper left corner of the login banner.

To upload an image to customize SOLIDserver

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Maintenance**, click on **Local files listing**. The page **Local files listing** opens.
3. Under the menu, tick the bullet **Custom images**. The sub-page **Custom images** opens.
4. In the menu, select . **Tools** > **Upload file**. The wizard **Import a file** opens, you can only upload one file at a time.
5. Click on **BROWSE** to look for the image of your choice on your computer.
6. In this new window, find the image you want to upload and select it.
7. Click on **Open**. The window closes and the wizard is visible again. In the field **File name**, the name of the selected image is displayed.
8. Click on **OK** to complete the operation. The report wizard opens and closes. The image is listed.

You can upload as many images as you need to the page *Local files listing*, follow the procedure above for each one of them.

Displaying an Image on SOLIDserver Login Window

Once you uploaded the image you want to display as new logo on the login page, you have to set its name as the value of the appropriate entry of the registry database.

To display an image in the appliance login page window

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. Filter the column **Name** with the keyword *logo.generic*. The list displays the entry *www.display.login_page.logo.generic*.
4. In the column **Value**, click on *<empty>*. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in the name of the image you want to display (name.extension). If you have not uploaded it yet to the page *Custom images* of the *Local files listing*, refer to the procedure [To upload an image to customize SOLIDserver](#).
6. Click on to complete the operation. The specified image replaces the default SOLIDserver logo. To see it open SOLIDserver in a different browser or log out.

If you want to display a different image on the Login page, edit the value of the entry *www.display.login_page.logo.generic* with the full name of an image you already uploaded to the *Local files listing*.

Displaying an Image on SOLIDserver Login Page Background

Once you uploaded the image you want to display as background of the appliance login page to the page *Local files listing*, you have to set its name as the value of the appropriate entry of the registry database.

To display an image as the background of the appliance login page

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. Filter the column **Name** with the keyword *bg.generic*. The list displays the entry *www.display.login_page.bg.generic*.
4. In the column **Value**, click on *<empty>*. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, type in the name of the image you want to display (name.extension). If you have not uploaded it yet to the page *Custom images* of the *Local files listing*, refer to the procedure [To upload an image to customize SOLIDserver](#).
6. Click on to complete the operation. The specified image is displayed as background of SOLIDserver login page. To see it open SOLIDserver in a different browser or log out.

If you want to display a different image on the Login page, edit the value of the entry *www.display.login_page.bg.generic* with the full name of an image you already uploaded to the page *Local files listing*.

Removing Custom Images from SOLIDserver Login Page

The registry database also allows to remove the image from the login page. You must edit the relevant item and empty its value.

To remove an image from the appliance login page

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. Filter the column **Name** with the keyword *generic*. The list displays the entries *www.display.login_page.logo.generic* (the logo) and *www.display.login_page.bg.generic* (the background image).
4. In the column **Value** of the entry of your choice, click on *<image-name>*. The wizard **Registry database Edit a value** opens.
5. In the field **Value**, empty the field.
6. Click on to complete the operation. The image is removed from SOLIDserver login page. To make sure of it, open SOLIDserver in a different browser or log out.

Customizing SOLIDserver Home Page Welcome Banner

SOLIDserver home page contains a welcome banner that can be edited to suit your needs: you can customize it with an image, change the message or even hide the banner altogether.

Only users of the group *admin* can edit the welcome banner.

Editing SOLIDserver Welcome Banner Title

By default, SOLIDserver home page contains a welcome banner containing the message *Welcome to SOLIDserver*, you can edit it if you want.

To edit the welcome banner title

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Dashboards**. The page refreshes.
2. In the drop-down list, select **Main Dashboard**. The page refreshes.
3. In the right corner of the welcome banner, click on . The wizard **Editing the welcome banner** opens.
4. In the field **Title**, replace the current message with the one of your choice.
5. Click on to complete the operation. The report opens and closes. The Main dashboard refreshes, the new message is visible.

Displaying an Image on SOLIDserver Welcome Banner

It is possible to add an image next to the title in the welcome banner. The size of this image does not matter as it is automatically resized to fit in the welcome banner.

To display an image on the welcome banner

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Dashboards**. The page refreshes.
2. In the drop-down list, select **Main Dashboard**. The page refreshes.

3. In the right corner of the welcome banner, click on . The wizard **Editing the welcome banner** opens.
4. Click on **BROWSE** to look for the image of your choice on your computer.
5. In this new window, find the image you want to upload and select it.
6. Click on **Open**. The window closes and the wizard is visible again. In the field **File name**, the name of the selected image is displayed.
7. Click on **OK** to complete the operation. The report opens and closes. The Main dashboard refreshes, the image is visible on the appliance welcome banner.

If you want to display a different image in the welcome banner, follow the procedure and select another image.

Keep in mind that the selected image(s) is saved on the page Local files listing and listed on the sub-page *Custom images*. For more details, refer to the section [Managing Files from the Local Files Listing](#).

Removing the Image from SOLIDserver Welcome Banner

At any time, the image displayed in the welcome banner can be removed.

To remove the image displayed on the Welcome banner

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Dashboards**. The page refreshes.
2. In the drop-down list, select **Main Dashboard**. The page refreshes.
3. In the right corner of the welcome banner, click on . The wizard **Editing the welcome banner** opens.
4. Tick the box **Remove the image from the banner**.
5. Click on **OK** to complete the operation. The report opens and closes. The Main dashboard refreshes, the image is no longer visible.

Hiding SOLIDserver Welcome Banner

You can remove the welcome banner from the SOLIDserver home page. In fact you are hiding it, a dedicated entry in the registry database is added to hide the banner.

To hide the welcome banner

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Dashboards**. The page refreshes.
2. In the drop-down list, select **Main Dashboard**. The page refreshes.
3. In the right corner of the welcome banner, click on **x**. The wizard **Hiding the welcome banner** opens.
4. Click on **OK** to complete the operation. The report opens and closes. The Main dashboard refreshes, the banner is no longer visible.

Displaying SOLIDserver Welcome Banner Again

To display the welcome banner on the SOLIDserver home page again, you need to delete the entry of the registry database that hides the banner.

To display the welcome banner again

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. Filter the column **Name** with the keyword *panel.home*. The list displays the entry *panel.home.welcome.delete*.
4. Tick the entry *panel.home.welcome.delete*.
5. In the menu, click on  **Delete**. The wizard **Delete** opens.
6. Click on to complete the operation. The page refreshes, the item is no longer listed.
7. In the sidebar, click on **Quit administration** to leave the module *Administration*.
8. In the sidebar, click on  **Dashboards**. The page refreshes. The banner is visible again.

Customizing the Interface Names and Fields

The Administration module provides a page dedicated to administrators that allows them to customize the interface labels, that is the default name of some fields and menus. To be precise, this page allows you to rename: fields name, menus name and pages and columns title. From *Language editor*, you can add entries to replace existing labels in the GUI.

There are two exceptions:

- You cannot edit the title of the Language editor page itself.
- You cannot rename the SOLIDserver home page welcome banner title using this page. For more details, refer to the section [Customizing SOLIDserver Main Dashboard Welcome Banner](#).

Besides, keep in mind that **the interface label customization applies to the language you chose to manage SOLIDserver with**. The label of the *English* interface field that you add to Language editor with a *Spanish* new name, is not edited.

To add a customized label

1. From any page or wizard within SOLIDserver, copy the name of a field, page, column or menu that you want to replace with your label.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Language editor**. The page **Language editor** opens.
4. In the menu, click on  **Add**. The wizard opens.
5. In the field **Key**, paste the value you want to replace. We recommend that you copy/paste the label name because Language editor is case sensitive.
6. If your appliance is displayed in English, in the field **English**, type in the new label value.
7. Click on to complete the operation. The entry is listed. Go back to the page where you copied the label to see the new name.

To delete a customized label

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Language editor**. The page **Language editor** opens.
3. In the column **Key**, click on the label name. The wizard opens.
4. Empty all the fields.
5. Click on **OK** to complete the operation. The entry is no longer listed. Go to the page the label is displayed on: it now displays the standard label.

To edit a customized label

1. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Language editor**. The page **Language editor** opens.
3. In the column **Key**, click on the label name. The wizard opens.
 - a. In the field **Key**, you can edit the label itself. This edits a different field, column or page, or nothing at all if it does not correspond to anything in the GUI.
 - b. If your appliance is displayed in English, in the field **English**, you can edit the label.
4. Click on **OK** to complete the operation. The entry is listed. Go to the page the label is displayed on: it now displays the edited label in the corresponding language.

Chapter 95. Managing Smart Folders

Smart folders is one the many customization opportunities provided by SOLIDserver. They allow to organize your items differently than on the listing pages. Smart Folders can be created, deleted, edited and/or shared with other users.

A smart folder is a database view that helps you organize data into a tree-like hierarchy. You can make a smart folder out of any list of items within the IPAM, DHCP, DNS, NetChange and Device Manager modules. This organization can have as many levels as you need and is composed of columns and/or meta-data (class parameters). Remember that this display is completely virtual and does not affect in any way your data.

Like the gadgets, smart folders can either be personal or shared with other users. For more details, refer to the section [Sharing Smart Folders](#).

Browsing Smart Folders

All the smart folders can be displayed in the side panel, next to the sidebar, from any module except Administration.

To manage them, you must go to the page My Smart Folders.

Displaying the Smart Folders in the Sidebar

To display the smart folders

1. In the sidebar, click on  **My Smart Folders**. The page refreshes, the side panel **Smart Folders** opens.
2. Click on  to display the content of a smart folder. Click on  to collapse a level or smart folder.
3. You can expand the side panel to the right or refresh the data displayed using the button **C**.

The Smart folders are displayed as follows:

- The icon  indicates an existing smart folder or the levels of hierarchy with a smart folder.
- Within a smart folder
 - Each level of hierarchy is indented.
 - Each level of hierarchy has a name followed by a number between brackets. This number indicates how many sub-levels or resources it contains.
 - The lowest level of the hierarchy is preceded by the icon of the resource where the smart folder was created: , , ...

Note that:

- If a level does not appear when you expand the hierarchy of the smart folder, it means there is no data to display.
- If a level is called *N/A*, it represents a level not relevant to the resource at the lowest level. For instance, if you create a smart folder from the page All addresses and include a level "pool name", if some IP addresses are not managed by a pool, they are listed under *N/A*.

Browsing the Smart Folders Database

From the page **My Smart Folders** you can manage edit, delete or share your smart folders.

To access the page **My Smart Folders**

1. From any page, in the top bar, select  **My Account > My Smart Folders**. The page **My Smart Folders** opens.
2. The columns on the page include all the information regarding each smart folder. They do not have a properties page.

Table 95.1. The columns on the page My Smart Folders

Column	Description
My Smart Folders	All the smart folders created.
All users	The smart folder visibility. <i>Yes</i> indicates that the smart folder is shared with the other users. <i>No</i> indicates that only the <i>User</i> who created it can see it.
User	The name of the user who created the smart folder(s) listed.
Type	The page where the smart folder was created, listed as follows: <code><module>:<page></code> .
Hierarchy	The smart folder hierarchy.

Adding Smart Folders

Smart Folders can be added, i.e. created, from any listing page within the modules IPAM, DHCP, DNS, NetChange and Device Manager. They allow to organize into a customized hierarchy the data displayed. Do not hesitate to filter the data at your convenience to visualize a tree displaying only the pieces of information of your choice.

To add a smart folder

1. Go to the page of your choice.
2. Filter the data if needed.
3. In the menu, select  **Alerts, gadgets & Smart Folders > Add a Smart Folder**. The wizard **Add a Smart Folder** opens.
4. In the field **Smart Folder Name**, name your Smart Folder.
5. Select the content of your smart folder:
 - a. In the drop-down list **Hierarchy**, select a column or a class parameter.
 - b. Click on . The selected value is moved to the list **Selected items in the hierarchy**.
 - c. Repeat these steps for as many columns and class parameters as needed.
6. Organize the hierarchy of your smart folder. The order impacts the final display of the Smart Folder.
 - a. In the list **Selected items in the hierarchy**, select a value and move it using the buttons  and .
 - b. To remove a value, select it in the list and click on . It moved back to the list *Hierarchy*.

7. Tick the box **Visible to the other users** if you want to share your Smart Folder with the other users.
8. Click on to complete the operation. The report opens and closes. The page is visible again.

Once created, the smart folder is listed in the Tree view. If you do not see it use the **C** button.

Editing Smart Folders

Smart folders can be edited at any time from the smart folder list. You can edit every column through this procedure, except the User column.

To edit a smart folder

1. From any page, in the top bar, select **My Account > My Smart Folders**. The page **My Smart Folders** opens.
2. Right-click over the name of the smart folder you want to edit. The contextual menu opens.
3. Click on . The wizard **Edit a Smart Folder** opens.
4. In the field **Smart Folder Name**, change the name if need be.
5. Add more columns and/or class parameters if need be:
 - a. In the drop-down list **Hierarchy**, select a column or a class parameter.
 - b. Click on . The selected value is moved to the list **Selected items in the hierarchy**.
 - c. Repeat these steps for as many columns and class parameters as needed.
6. Edit the hierarchy of your smart folder if need be:
 - a. In the list **Selected items in the hierarchy**, select a value and move it using the buttons and .
 - b. To remove a value, select it in the list and click on . It moved back to the list *Hierarchy*.
7. Tick the box **Visible to the other users** if you want to share your Smart Folder with the other users. You can untick it as well.
8. Click on to complete the operation. The report opens and closes. The page is visible again. Your changes are displayed on the page *My Smart Folders* and in the *Tree view*.

Once edited, the smart folder new configuration can be displayed in the Tree view. Click on **C** to refresh the display.

Sharing Smart Folders

For each smart folder created you have the possibility to share it with other users or make it accessible only to you. There are two ways to do so, you can either choose to share the smart folder when creating it - for more details, refer to the part [Adding Smart Folders](#) -, or you can choose to share or hide it from the smart folders list.

To make a smart folder visible to all users

1. From any page, in the top bar, select  **My Account** > **My Smart Folders**. The page **My Smart Folders** opens.
2. In the list, tick the Smart folder you want to share. Filter the Smart Folders if need be.
3. In the menu, select  **Edit** > **Visible to all users** > **Yes**. The wizard **Smart Folder visibility** opens.
4. Click on  to complete the operation. The report opens and closes. The page **My Smart Folders** is visible again. In the column **All Users**, the Smart Folder is marked **Yes**.

The same procedure allows you to make a Smart Folder visible only to you.

To make a smart folder visible only to you

1. From any page, in the top bar, select  **My Account** > **My Smart Folders**. The page **My Smart Folders** opens.
2. In the list, tick the Smart folder you want to make visible only to you. Filter the Smart Folders if need be.
3. In the menu, select  **Edit** > **Visible to all users** > **No**. The wizard **Smart Folder visibility** opens.
4. Click on  to complete the operation. The report opens and closes. The page **My Smart Folders** is visible again. In the column **All Users**, the Smart Folder is marked **No**.

Deleting Smart Folders

Smart Folders can be deleted from the smart folder list. You can delete one or several smart folders at a time.

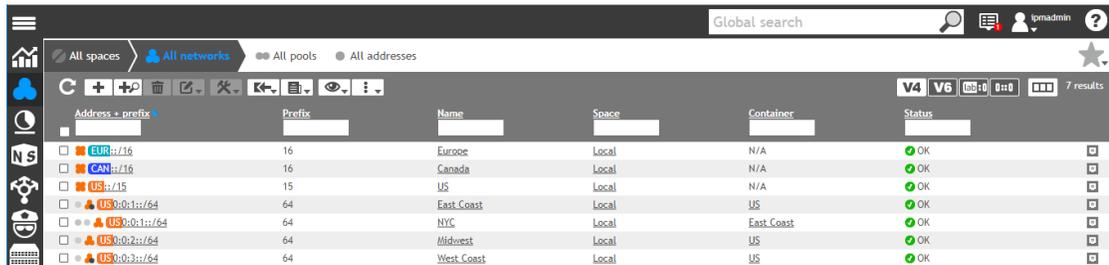
To delete a smart folder

1. From any page, in the top bar, select  **My Account** > **My Smart Folders**. The page **My Smart Folders** opens.
2. In the list, select the Smart Folder you want to delete. Filter the Smart Folders if need be.
3. In the menu, click on  **Delete**. The wizard **Delete** opens.
4. Click on  to complete the operation. The report opens and closes. The page **My Smart Folders** is visible again. The selected Smart Folder(s) is no longer listed.

Chapter 96. Managing IPv6 Labels

Labels provide a visual aid for IPv6 addresses management that allows to display the letters and colors of your choice above a defined part of the addresses.

They allow to gather at a glance IP addresses belonging to a common container in the modules **IPAM**, **DHCP**, **Application** and **NetChange**.



Address + prefix	Prefix	Name	Space	Container	Status
<input type="checkbox"/> EUR::/16	16	Europe	Local	N/A	OK
<input type="checkbox"/> CAN::/16	16	Canada	Local	N/A	OK
<input type="checkbox"/> US::/15	15	US	Local	N/A	OK
<input type="checkbox"/> US:0:0:1::/64	64	East Coast	Local	US	OK
<input type="checkbox"/> US:0:0:1::/64	64	NYC	Local	East Coast	OK
<input type="checkbox"/> US:0:0:2::/64	64	Midwest	Local	US	OK
<input type="checkbox"/> US:0:0:3::/64	64	West Coast	Local	US	OK

Figure 96.1. Example of a geographical distribution of labels in the IPAM

In the example above, the labels are named after the block-type and subnet-type networks. The colors reflect the hierarchy. Also, the label goes above, and therefore hides, the configured address, whether it is a full IP address or part of an address.

Limitations

- When configuring a label, you must type in the uncompressed version of the IP address. For instance, to create a label for a network starting with the address `12::`, you must specify `0012::`.
- A label applies to IPv6 addressing regardless of the module, once set it applies to IPAM, DHCP, Application and/or NetChange. Within the IPAM, if you have common network start addresses among several spaces or networks, they all have the same label (see the East Coast and NYC network labels in the example above).

Adding Labels

The labels are all managed from the same wizard, accessible in the menu **⋮** **Extra options**. You can add as many labels as you need on the following pages:

- In the module IPAM, you can manage labels on the IPv6 pages *All networks*, *All pools* and *All addresses*.
- In the module DHCP, you can manage labels on the IPv6 pages *All scopes*, *All ranges*, *All leases* and *All statics*.
- In the module Application, you can manage the labels on the page *All nodes*.
- In the module NetChange, you can manage the labels on the IPv6 pages *All addresses* and *All routes*.

To add a label

1. Depending on your needs, in the sidebar:
 - a. Go to **IPAM** > **Networks**, **Pools** or **Addresses**. The page opens.

- b. Go to  **DHCP > Scopes, Ranges, Leases or Statics**. The page opens.
 - c. Go to  **Application > Nodes**. The page opens.
 - d. Go to  **NetChange > Addresses or Routes**. The page opens.
2. If you accessed IPAM, DHCP or NetChange, on the right-end side of the menu, click on . The page refreshes and the button turns black.
 3. In the menu, select  **Extra options > Configure IPv6 labels**. The wizard **Configure IPv6 labels** opens.
 4. In the field **IPv6**, type in or paste the uncompressed address you want to label, or part of it.
 5. In the field **Label Name**, type in the label name of maximum 3 characters, letters or numbers.

The label is visible in the *Preview*.
 6. Under **Choose a color**, click on the color of your choice.

The color is visible in the *Preview*.
 7. Click on . The label is moved to the **List of labels**. Repeat these steps for as many labels as you need.
 8. Click on  to complete the operation. The list is visible again.

Once created, the labels need to be displayed manually, as detailed below.

Displaying or Hiding Labels

Once you created labels, you can choose to display or hide them.

To display/hide the labels

1. Depending on your needs, in the sidebar:
 - a. Go to  **IPAM > Networks, Pools or Addresses**. The page opens.
 - b. Go to  **DHCP > Scopes, Ranges, Leases or Statics**. The page opens.
 - c. Go to  **Application > Nodes**. The page opens.
 - d. Go to  **NetChange > Addresses or Routes**. The page opens.
2. If you accessed IPAM, DHCP or NetChange, on the right-end side of the menu, click on . The page refreshes and the button turns black.
3. On the right-end side of the menu, click on  **Use IPv6 labels**. The configured labels are visible and the button turns black.
4. On the right-end side of the menu, click on  **Do not use IPv6 labels**. The configured labels are no longer displayed and the button turns white.

Editing Labels

You can edit existing labels directly in their creation wizard.

To edit a label

1. Depending on your needs, in the sidebar:
 - a. Go to  **IPAM** > **Networks**, **Pools** or **Addresses**. The page opens.
 - b. Go to  **DHCP** > **Scopes**, **Ranges**, **Leases** or **Statics**. The page opens.
 - c. Go to  **Application** > **Nodes**. The page opens.
 - d. Go to  **NetChange** > **Addresses** or **Routes**. The page opens.
2. If you accessed IPAM, DHCP or NetChange, on the right-end side of the menu, click on . The page refreshes and the button turns black.
3. In the menu, select  > **Extra options** > **Configure IPv6 labels**. The wizard **Configure IPv6 labels** opens.
4. In the field **List of labels**, select the label you want to edit.
5. Edit the label **IPv6**, **Label Name** and/or **Color**.
6. Click on  to save the changes. The label is no longer listed in the field. Repeat these steps for as many labels as you need.
7. Click on  to complete the operation. The list is visible again.

Keep in mind that the labels need to be displayed manually. For more details, refer to the section [Displaying or Hiding Labels](#).

Deleting Labels

Like for the edition, you can delete existing labels directly in their creation wizard.

To delete a label

1. Depending on your needs, in the sidebar:
 - a. Go to  **IPAM** > **Networks**, **Pools** or **Addresses**. The page opens.
 - b. Go to  **DHCP** > **Scopes**, **Ranges**, **Leases** or **Statics**. The page opens.
 - c. Go to  **Application** > **Nodes**. The page opens.
 - d. Go to  **NetChange** > **Addresses** or **Routes**. The page opens.
2. If you accessed IPAM, DHCP or NetChange, on the right-end side of the menu, click on . The page refreshes and the button turns black.
3. In the menu, select  > **Extra options** > **Configure IPv6 labels**. The wizard **Configure IPv6 labels** opens.
4. In the field **List of labels**, select the label you want to delete. You can only delete labels one at a time.
5. Click on . The label is no longer listed in the field. Repeat these steps for as many labels as you need.
6. Click on  to complete the operation. The list is visible again.

Chapter 97. Configuring Classes

Classes allow to create properties and behaviors that apply only to the resources of your choice, they allow to tailor databases to their needs.

From the page **Class Studio**, users of the group *admin* can create, edit, rename and duplicate classes: they can also move them from one directory to the other or even from one resource to the other.

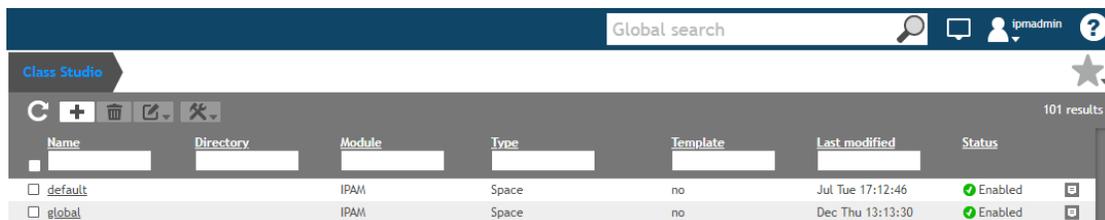


Figure 97.1. The page Class Studio

Before configuring classes, keep in mind that:

1. **The class is a container.** All classes are listed on the page Class Studio. For more details, refer to the section [Managing Classes](#).
2. **The class's content is managed independently.** All the parameters of a class, its class objects, are configured in a dedicated wizard. For more details, refer to the section [Configuring the Classes' Content](#).

When a class is selected, its parameters are available for configuration in the addition/edition wizard of the resource. If the class is created empty, the wizard is unchanged.

Class parameters can be inherited and/or propagated from one module level to another. For more details, refer to the section [Inheritance and Propagation](#).

Within SOLIDserver, you can apply existing classes to resources from their addition/edition wizard.

There are 3 types of classes:

- **Default** is a class associated by default with every type of resource. These classes are automatically applied to a resource in the addition and edition wizard and set the object's advanced properties. For more details, refer to the chapter [Managing Advanced Properties](#).

You cannot delete or edit *default* classes.

- **Global** is a class associated by default with every type of resources as well. These classes are automatically applied to a resource in the addition and edition wizard and set the resource's **Meta-Data**, the user defined fields of your choice.

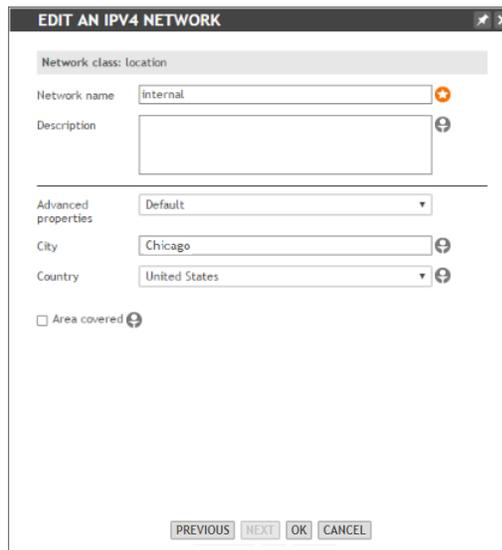
You can edit but not delete *global* classes. For more details, refer to the section [Editing Classes](#).

- **Customized classes** are all the classes you create from the page Class Studio. These classes must be manually selected and applied to resources in their addition and edition wizard. The wizard automatically detects that a custom class is enabled for the type of resource and provides, before anything else, the list **<resource> class**, where you select the class of your choice or *None*.

You can manage (add, edit, delete...) customized classes. For more details, refer to the section [Managing Classes](#).

Note that customized classes can be set as resource of groups of users. If users have management rights over objects configured with a customized class but the class is not among their resources, they cannot display or edit the class parameters and they clear the value of all the class parameters when they edit a resource. For more details regarding user resources, refer to the section [Assigning Resources to a Group](#).

Applying an enabled class may change the fields available in the addition/edition wizard of the resource. In the example below, an administrator configured a network class called *location* including an input field labeled *City*. When editing the network *internal*, a user selected this class, and can now specify that the network is located in the *City* of *Chicago*.



The screenshot shows a window titled "EDIT AN IPV4 NETWORK". At the top, it says "Network class: location". Below that, there are several input fields: "Network name" with the value "internal", "Description" (empty), "Advanced properties" (set to "Default"), "City" (set to "Chicago"), and "Country" (set to "United States"). There is also an unchecked checkbox labeled "Area covered". At the bottom of the window, there are four buttons: "PREVIOUS", "NEXT", "OK", and "CANCEL".

Figure 97.2. Example of an input class object labeled "City"

Note that applied classes can also be used to set automatic templates to display specific columns on a page based on the class applied to the container of the resources listed. For more details refer to the procedure [To add an automatic template](#).

Browsing Class Studio Database

In the module Administration, the page Class Studio allows to display all existing classes and their content.

The classes are all listed on the page, while their content must be listed and managed from a dedicated wizard.

Browsing Classes

To display the list of classes

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

- In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.

By default, Class Studio displays as many global and default classes as there are resources within SOLIDserver. The page columns are described in the table below:

Table 97.1. Class Studio columns

Column	Description
Name	The class name. It might be set to <i>default</i> , <i>global</i> , or defined by the user in the case of customized classes. Only customized classes can be renamed.
Directory	The directory in which the class is located. A directory can only be created upon addition of a class. Still, a class can be moved from one directory to another at any given time.
Module	The module of the resource for which the class is set. A class can only be moved from a module to another when not in use.
Type	The type of resource for which the class is set (DHCP groups, DNS servers, etc.). A class can only be moved from a type of resource to another when not in use.
Template	The class template configuration: it is either a template or not. Its value can be <i>yes</i> or <i>no</i> . For more details, refer to the chapter Managing IPAM Templates .
Last modified	The time and date of the last modification made on the class.
Status	The class activation: <i>Enabled</i> or <i>Disabled</i> . When disabled, a class is neither applied nor listed among the class available when adding or editing the resource it could apply to.

To display a class properties page

Only users of the group *admin* can perform this operation.

- In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
- At the end of the line of the class of your choice, click on **⌘**. The properties page opens.

Browsing Class Objects

The classes configuration is available in a pop-up window that allows to add and edit the class objects of the class. **To edit classes, your browser must allow pop-up windows.**

This pop-up window, named *Class Editor*, opens when you click any class name on the page Class Studio. It is divided vertically to display: on the left a **creation panel** and on the right **the list of class objects**. You can sort these objects using the drop-down list.

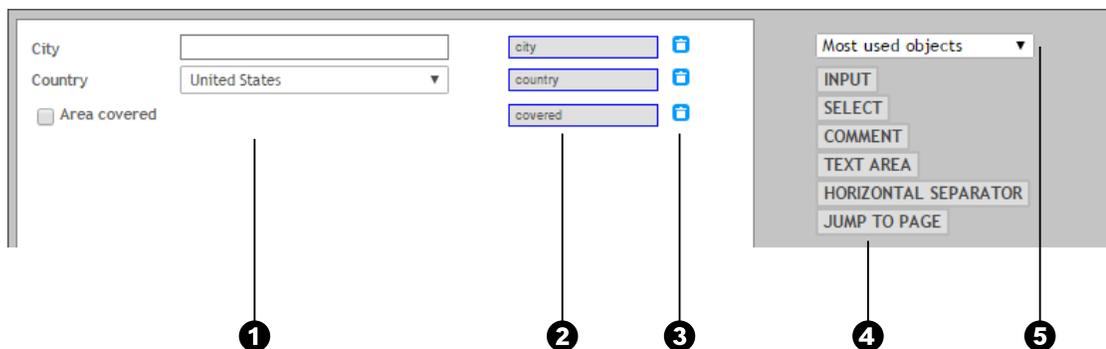


Figure 97.3. The pop-up window Class Editor

- ❶ This section of the window displays the class objects' *Label* as it appears in the wizard.
- ❷ The gray fields display the *Name* of the class object.
- ❸ The  icon allows to delete the class objects.
- ❹ The gray section of class editor displays the available class objects.
- ❺ The drop-down list allows to filter the class object types.

To open Class Editor from Class Studio

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the column **Name**, click on the class you want to edit. The pop-up window **Class Editor (<hostname>/<user>)** opens. The class name is displayed at the end of the URL field as such: `<class-name>.class` .

Class Editor allows to edit any of your own classes and the class *global* of any resource.

Managing Classes

From Class Studio, you can add, edit, rename, duplicate, move, stop using and/or delete classes. Once created you can apply classes to the type of resource they were configured for.

Keep in mind that two types of class are available by default but their management is more limited:

- **Default classes:** you cannot edit or delete any *default* class. However, you can choose the advanced properties you want to use, for more details refer to the chapter [Managing Advanced Properties](#).
- **Global classes:** you can edit each *global* class from the menu  **Extra options > Meta-data** but you cannot delete them. For more details, refer to the section [Editing Classes](#) below.

In this section, we only discuss the classes themselves, to configure their content, i.e. the class objects that tailor them to your needs, refer to the section [Configuring the Classes' Content](#).

You can create and manage classes for the following resources, the columns correspond to the drop-down lists available in the class addition wizard.

Table 97.2. Resources that can be configured with a class

Module	Type
Administration	SOLIDserver (appliance).
Application	Application.
DHCP	Servers, scopes, ranges, groups and statics both in IPv4 and IPv6.
DNS	Servers, views and zones.
Device Manager	Devices and Ports & interfaces.
IPAM	Spaces, networks, pools and addresses both in IPv4 and IPv6.
NetChange	Network devices and ports.
Rights & delegation	Groups and users. <i>The groups of users are managed from the module Administration.</i>
SPX	Autnums. <i>The Autnums are managed from the module IPAM.</i>

Module	Type
VLAN manager	VLAN domains and ranges.
VRF	VRFs.
Workflow	Requests. <i>The resource Requests only applies to outgoing requests.</i>

Adding Classes

You can create specific customized classes to be applied individually and manually to your resources when you add or edit them.

Note that you can also edit the *global* classes but this would automatically affect all the resources the class is set for. For more details, refer to the section [Editing Classes](#).

To add a class

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. In the menu, click on **Add**. The wizard **Add a new class** wizard.
4. Fill in the following fields:

Table 97.3. Class addition parameters

Field	Description
Filename	Name your class. The name cannot contain any special characters. This field is mandatory.
Sub directory	You can fill in the directory where you want to save your class. If it does not exist, it is created. On the wizards class selection page, classes placed in a directory are displayed as such: <code><directory>/<class></code> . This field is optional.
Module	Select the module of your choice. This field is mandatory.
Type	Select the resource of your choice. This field is mandatory.
Enable class	Tick this box if you want to enable the class upon creation. If a class is not enabled, it is ignored in the database and not listed in the addition/edition wizard of the resource it applies to. This field is optional as you can enable it later on, For more details, refer to the section Enabling or Disabling Classes .

5. Click on to complete the operation. The report opens and closes. The class is listed.

Once you added a class, keep in mind that:

- **A class is empty by default**, whether it is a *global* or a customized one. Once created, you can click on the name of a class name to add and configure its class objects through Class Editor. For more details, refer to the section [Configuring the Classes' Content](#).
- **A class must be enabled to be used**. If a customized class is not enabled, it is not available in the addition and edition wizards of the resource. For more details, refer to the section [Applying Classes](#).

Only *global* classes are enabled by default and can be edited and automatically integrated to the wizards of the resources they are set for.

Applying Classes

Once at least one customized class is created and enabled, the addition and edition wizard of the resource it can be applied to displays a dedicated page `<resource> class`.



Figure 97.4. The page Network class in an addition/editing wizard

This page allows to select and apply a class, that is to say load and configure its class objects in the wizard.

Note that in some modules, you can configure and apply classes at many levels. When you are adding resources at low levels without filtering the page to display the content of a specific container, you need to select the resource container(s). If classes are applied at container levels, you need to select a `<resource> class` for each container level, this does not load class objects in the wizard, it allows to filter the list of potential containers for the resource you are adding.

To apply a class

Only users of the group `admin` can perform this operation.

1. Go to the page `All <resources>` of a resource for which you enabled a class.
2. Add or edit a resource. The wizard opens.
3. On the page `<resource> class`, select a value.
 - a. If no class was added at higher level:

Table 97.4. Class selection options

Option	Description
<code>None</code>	Select this option if you do not want to apply any customized class to the resource.
<code><class-name></code>	Select a class to apply it on the resource and load its content in the wizard. Classes belonging to a directory are listed as follows: <code><directory>/<class-name></code> .

- b. If you need to select a container and at least one class was applied at higher level:

Table 97.5. Container class selection options

Option	Description
<i>All</i>	Select this option to display the list of all the available containers on the next page, whether they are configured with or without a customized class.
<class-name>	Select a class to display the list of all the containers using it on the next page. Classes belonging to a directory are listed as follows: <directory>/<class-name>.
<i>No class</i>	Select this option to only display the containers configured without customized class on the next page. If the option is not available, it means that all the containers are configured with a class.

- Click on **NEXT** to load the next page. The class object fields are displayed with the standard addition/edition fields, you can configure them.

If you did not filter the list and are adding a low level resource, you need to select a container and click on *NEXT* and repeat these actions until you get to the page <resource> class of the resource you want to add.

- Click on **OK** to complete the operation. The report opens and closes.

Editing Classes

You can edit the content of a customized class at any time, that is to say edit the class objects it contains or edit their order.

To edit a class

Only users of the group *admin* can perform this operation.

- Make sure your browser allows pop-up windows.
- In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
- In the column **Name**, click on the class you want to edit. **Class Editor** opens.
- Edit the class object according to your needs following the procedure that suits your needs in the section [Adding Class Objects](#).
- Click on **OK** to complete the operation. The object is updated in the creation panel.

Keep in mind that you can also edit the configuration of the class *global* of a specific resource directly from its management page *All <resources>*, except for the following objects:

- DHCP groups, leases, ACLs, ACL entries, option definitions and failover channels.
- DNS views, DNSSEC keys, RRs, RPZ zones and RPZ rules.
- NetChange configurations, routes, addresses, VLANs and discovered items.
- VLAN Manager VLANs.
- VRF Route Targets.
- SPX policies.

To edit the class global of a resource from its management page

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. Click on of the resource of your choice. The page opens.
3. In the menu, select **⋮** **Extra options** > **Meta-data**. Class Editor opens and displays the class objects of the class *global* of the chosen resource.
4. Edit the class object according to your needs following the procedure that suits your needs in the section [Adding Class Objects](#).
5. Click on **OK** to complete the operation. The object is updated in the creation panel.

Renaming Classes

A customized class can be renamed at any time from its properties page. Renaming a class does not affect the class objects it contains. Once a class has been renamed, it is updated on the properties page of the concerned resources.

To rename a class

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. At the end of the line of the class of your choice, click on **⌵**. The properties page opens.
5. In the menu, select **✎ Edit** > **Rename**. The wizard **Rename class** opens.
6. In the field **Old**, the current class name is displayed.
7. In the field **New Name**, type in the new name for the class.
8. Click on **OK** to complete the operation. The class new name is displayed in the panel and modified in the list.

Duplicating Classes

You can duplicate customized classes. These duplicates can then be edited and renamed to manage them more easily, for instance you might need to apply them to other types of resource or even move them.

Duplicating classes can be useful since object values set for a resource are automatically inherited by the resources it contains. For instance, if the value "*Chicago*" is set for a block-type network through an input field "*city*", it is automatically inherited by the subnet-type networks it contains if said subnet-type network also has an input field named "*city*".

To duplicate a class

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. Tick the class(es) you want to duplicate.
4. In the menu, select **✎ Edit** > **Duplicate**. The wizard **Duplicate class** opens.

5. Click on **OK** to complete the operation. The duplicated class is listed and named as such: *copy_<original class name>*.

Moving Classes

In contrast with the *default* and *global* classes that are hard linked to the resources they are set for, customized classes can be moved from a directory to another or even from a type of resource to another. For instance, a class created for DNS servers can be moved and made available for a completely different type of resource, like the DHCP ranges.

To move a class

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. Tick the class(es) you want to move.
4. In the menu, select **Edit > Move**. The wizard **Move class** wizard.
5. In the field **Sub-directory**, type in a directory if need be. It can be a new directory for the class or an existing one.
6. In the drop-down list **Module**, select a module for the class. It can be the same one or a new one.
7. In the drop-down list **Type**, select a resource to which the class should be applied. It can be the same one or a new one.
8. Click on **OK** to complete the operation. The report opens and closes. The data is updated in the list.

Changing or Stop Using Classes

At any time, you can decide not to use a particular class on the resource of your choice. For instance, you might decide not to use a class that you want to delete or need to use another class for a particular resource.

As classes must not be used at all in SOLIDserver to be deleted, the following procedure might come in handy. Keep in mind that the columns layout on the page can help you find the resources using a class. For more details, refer to the section [Customizing the List Layout](#).

To change or stop using a class on a specific resource

Only users of the group *admin* can perform this operation.

1. Go to the page *All <resources>* of a resource for which you enabled a class.
2. Click on of the resource of your choice. The page opens.
3. At the end of the line of the resource of your choice, click on **⌵**. The properties page opens.
4. In the panel **Main properties**, click on **EDIT**. The related edition wizard opens.
5. Click on **NEXT** until you reach the page **<Resource> class** of the wizard.
6. In the list **<Resource> class**, select *None* or a class different from the one you intend to delete.
7. Click on **NEXT** until you reach the last page of the wizard.

8. Click on to complete the operation. The report opens and closes. The class has been dissociated from the resource.

Enabling or Disabling Classes

Upon addition, a customized class can either be enabled straight away or left disabled. Since deleting classes may result in unwanted complications, disabling classes allows to store them for future use rather than deleting them.

Keep in mind that the *default* and *global* classes cannot be disabled and are automatically applied on the resources they are set for.

There are two ways to enable a class, either from Class Studio or from a listing page.

To enable/disable a class from class studio

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. Tick the class(es) of your choice.
4. In the menu, select  **Edit** > **Enable class** or **Disable class**. The wizard opens.
5. Click on to complete the operation. The report opens and closes, the page refreshes. The class is marked as  *Enabled* or  *Disabled* in the column **Status**.

To enable/disable a class from a listing page

Only users of the group *admin* can perform this operation.

1. Go to the listing page of your choice.
2. In the menu, select  **Extra options** > **Classes configuration**. The wizard **<object> Classes Configuration** opens.
3. In the list **Classes library**, select a class and click on  to enable it. The group is moved to the list **Enabled classes**. Repeat these actions for as many classes as needed.
4. In the list **Enabled classes**, the classes enabled for this type of object are listed. You can remove one (or several) classes from that list if you want to disable them: select the group and click on . The class is listed back in the list **Classes library**.
5. Click on to complete the operation. The report opens and closes, the listing page is visible again.

Deleting Classes

Only customized classes can be deleted. Keep in mind that:

- You can delete customized classes if and only if they are not used by any resource within SOLIDserver. Therefore, you might need to stop using the class before deleting it. For more details, refer to the section [Changing or Stop Using Classes](#).
- Deleting a class deletes the class objects it contained and displayed on the resources properties page. You might simply want [To enable/disable a class from class studio](#) or [To enable/disable a class from a listing page](#) to use it again later.

To delete a class

Only users of the group *admin* can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
3. Tick the class(es) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete class**.
5. Click on to complete the operation. The report opens and closes. The class has been deleted is no longer listed.

Configuring the Classes Content

Users of the group *admin* can configure the classes, in other words set the class objects that define what the class does when it is applied to an object. Once set, a class customizes the resource's addition and edition wizards with extra pages, comments, boxes, lists, input fields... that can be prefilled with a value retrieved automatically or set manually.

As every new class is empty, it does not bring any changes to the resource it applies to. You need to add class objects within a class to configure its behavior. They can be added to *global* and customized classes via Class Editor. For more details, refer to the section [Browsing the Class Objects](#). Keep in mind that:

- Customized classes only affect the addition/edition wizard of a resource if they are enabled. Whereas *global* classes automatically affect them.
- The object values set for a resource are automatically inherited by the objects it contains. For instance, if the value *Chicago* is set for a block-type network through an input field *city*, it is automatically inherited by the subnet-type networks it contains if said subnet-type network also has an input field named *city*.

You can tick several objects to set the inheritance or propagation property of their class parameters. For more details, refer to the section [Inheritance and Propagation](#).

- To configure classes, **your browser must allow pop-up windows**.

For each class, Class Editor includes a large library of **class objects** (formerly WDOM objects) gathered in groups. These groups can be selected one by one in the drop-down list:

Most used objects

Contains the most frequently used class objects, no matter what module or resources the classes can apply to:

Class object	Description
Input	Displays an input field that allows users to add data on one line.
Select	Displays a drop-down list that allows users to add data from a list of manually set or automatically retrieved values. These values can come from a CSV file, a service list or a custom DB.
Comment	Displays a <i>Notice</i> , <i>Warning</i> or <i>Information</i> message that contain the text of your choice.
Text area	Displays a large input field that allows users to add data on several lines, it can contain up to 3600 characters.

Configuring Classes

Class object	Description
Horizontal separator	Displays a colored line, either red, green or blue, that allows to separate and organize the class fields according to your needs.
Jump to page	Splits the wizard in several pages, it therefore adds a button <i>NEXT</i> at the bottom of the page.

IP address management

Contains class objects dedicated to IPAM classes:

Class object	Description
Hide IP alias	Allows to hide the alias request page when assigning an IP address. For more details, refer to the section Configuring and Managing IP Address Aliases .
Force prefix	Allows to force a prefix on a network.

DHCP management

Contains class objects dedicated to DHCP classes:

Class object	Description
Select DHCP server	Displays a drop-down list containing all the existing DHCP servers.
Select DHCP scope	Displays a drop-down list containing all the existing DHCP scopes.
Select DHCP range	Displays a drop-down list containing all the existing DHCP ranges.
Select DHCP static	Displays a drop-down list containing all the existing DHCP statics.
DHCP options	Displays a large set of DHCP options. For more details, refer to the appendix DHCP Options .

DNS management

Contains class objects dedicated to DNS classes:

Class object	Description
Select DNS server	Displays a drop-down list containing all the existing DNS servers.
Select DNS zone	Displays a drop-down list containing all the existing DNS zones.
Select DNS domain	Displays a drop-down list containing all the existing DNS domains.

All Objects

Contains all the available class objects, including the ones describes above. The remaining class objects are the following ones:

Class object	Description
Select class	Displays a drop-down list containing the existing classes for a specific resource.
Autocompletion	Displays an input field that can be configured to provide suggestions or automatic data completion through predefined values.
Checkbox	Displays a customizable checkbox, or box.
DHCP shared network	Displays a drop-down list containing all the existing scopes used as shared networks.
Counter	Displays an incremental counter.
Force class	Forces a class on every resource of a container.
Force VLSM	Allows to force the VLSM status of all subnet-type networks created to non-terminal. This class object can be applied to spaces and networks (block-type and subnet-type).
Hidden data	Allows to hide one of your class objects when configuring a class. The value of the field is set for the resource but not displayed in the addition and edition wizards.

Class object	Description
Icon	Allows to associate an icon to a class. This icon appears next to the class name in the column <i>Class</i> .
Include class	Allows to embed another class objects to your class.
Multiple input	Displays a list under a field <i>Input</i> where you can store the values specified.
Multiple select	Displays a drop-down list and a text area to select and store multiple values at the same time. These can be <i>fixed values</i> or value automatically imported from a CSV file, a service list or a custom DB.
Object name	Sets some naming convention rules for a resource (E.g. <i>Network name = <country>-<network_number></i>).
Pre-defined variable	Predefined variables can be seen as full-fledged class objects with only one value and purpose.
Time stamp	Displays a field containing the date and time of the creation of the resource.
Upload file	Displays a button and a field that allows uploading a files stored on a local computer to the <i>tmp</i> folder of the appliance.
Owner	Displays a field containing the name of the user who created the resource.

Note that you can display on any listing page a column matching each of your customized classes to order and filter the list based on the classes or class parameters applied on certain resources. For more details, refer to the section [Customizing the List Layout](#).

Adding Class Objects

You can as many class object as you want to your classes. Note that any class object can be added to any class, but some are designed for specific resources.

Input

The class object *Input* allows to display an input field in the wizard that users can fill with a data string.

Note that an *Input* is different from a *Multiple input*. For more details, refer to the section [Multiple input](#).

To add an input field

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on [Input](#). The wizard **Input** opens.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.

8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the field **Input field maximum length**, type in the maximum number of characters, including spaces, that users can type in the field. By default, the maximum field length is *64*.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Not editable** to prevent users from editing the class object's value. If you tick the box, the object appears in gray in the addition and edition wizards.
 - b. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
 - c. In the field **Constructor**, you can type in the value of other class objects of the class and use them as variables to automatically overwrite the *Name* of the field with the data of your choice in the wizard.

For example, you could type in `%v{<value1>}, %v{<value2>}` where `<value#>` is the value of an existing class object in the class. If `<value1>` is a *city* and `<value2>` is a *state*, the field *Name* would be replaced with *Chicago, Illinois* in the wizard.

- d. In the drop-down list **Predefined format**, you can select a format for the *Name* to be valid. It can either be an *IP address (v4)*, *IP address (v6)*, *Text*, *Unsigned integer*, *Signed Integer*, *Domain name*, *FQDN Host*, *MAC address* or *Email address*.
- e. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.
- f. In the field **Regex match**, you can type in a regular expression that checks the syntax of the value specified in the field. For more details, refer to the section [Class Studio Syntax](#).
- g. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- h. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- i. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select

The class object *Select* allows to display a drop-down list in the wizard. The data listed can be:

- *Fixed values* added to the list directly from the wizard
- *CSV values* retrieved from a CSV file
- *Service list values* picked from the SOLIDserver services. All services and parameters are described in the API reference guide on our website ¹.
- *Custom DB values* retrieved directly from your custom databases. We strongly recommend using this type rather than a CSV file. For more details, refer to the chapter [Custom DB](#).

Note that a *Select* is different from a *Multiple select*. For more details, refer to the section [Multiple select](#).

To add a select drop-down list using fixed values

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on . The wizard **Select** wizard.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Select type**, select *Fixed values*. The wizard refreshes.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

¹At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

- a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
- b. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
- c. You can tick the box **Reload on change** if you want to reload the wizard page when a value is selected in the drop-down list.
- d. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- e. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- f. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

- g. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.

11. Click on **[NEXT]**. The last page of the wizard opens.
12. In the field **Key**, type in the object name as it should be saved in SOLIDserver database (string of characters: *_a-z0-9* only). To prevent GUI conflicts, avoid names that are already used in the code such as: *site*, *mac-addr*, *gateway*, *vlan*, *domain*, *user*, *port*, *password*... The field **Label/Key** autopopulates.
13. In the field **Label**, type in the word string, corresponding to the key, as it should be displayed in the list. The field **Label/Key** autopopulates following the format *<Label>#<Key>*.
14. Click on **[+]**. The value is moved to the list **Options**.

15. Repeat these actions for as many values as needed. You can use to remove one by one values from the list, or and to reorganize them.
16. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add a select drop-down list using a CSV file

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on . The wizard **Select** wizard.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Select type**, select *CSV values*. The wizard refreshes.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
 - b. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
 - c. You can tick the box **Reload on change** if you want to reload the wizard page when a value is selected in the drop-down list.
 - d. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- e. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- f. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

- g. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.
11. Click on **[NEXT]**. The last page of the wizard opens.
 12. In the field **CSV file**, type in the complete path of the file stored in the appliance.
 13. In the field **Value column**, type in the number of the column in the CSV file containing the values to retrieve.
 14. In the field **Label column**, type in the number of the column in the CSV file containing the labels corresponding to the values to retrieve.
 15. In the field **Filter column**, type in the number of the column used to match certain rows.
 16. Click on **[+]**. The value is moved to the list **Filter**. You can use **[-]** to remove one by one values from the list, or **[<]** and **[>]** to reorganize them.
 17. Click on **[OK]** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by **[★]** if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add a select drop-down list using service list values

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **⚙ Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on **[Select]**. The wizard **Select** wizard.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site*, *mac-addr*, *gateway*, *vlan*, *domain*, *user*, *port*, *password*...

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Select type**, select *Service list values* or *Manual*. The wizard refreshes.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

- a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
- b. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
- c. You can tick the box **Reload on change** if you want to reload the wizard page when a value is selected in the drop-down list.
- d. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- e. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- f. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

- g. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.

11. Click on **[NEXT]**. The last page of the wizard opens.
12. In the field **Services**, start typing in the name of the service to call. The matching services are listed, select the one that suits your needs.

13. In the field **Key**, type in the name of the input parameter corresponding to the values to retrieve.
14. In the field **Label**, type in the name of the input parameter corresponding to the labels associated to these values.
15. In the field **Where**, type in an SQL condition to filter the retrieved values if need be.
16. In the field **Order by**, type in an SQL condition to sort the results if need be.
17. In the field **Limit**, type in the maximum number of results to display.
18. In the field **Tags**, type in an SQL condition to filter the retrieved class parameters if need be. You might need assistance from Efficient IP support team to fill in this field.
19. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add a select drop-down list using Custom DB values

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on . The wizard **Select** wizard.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Select type**, select *Custom DB*. The wizard refreshes.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
 - b. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
 - c. You can tick the box **Reload on change** if you want to reload the wizard page when a value is selected in the drop-down list.
 - d. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- e. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- f. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

- g. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.

11. Click on . The last page of the wizard opens.
12. In the field **Custom DB name**, type in the name of the Custom DB of your choice. For more details on SOLIDserver Custom DB, refer to the chapter [Custom DB](#). The field autocompletes.
13. In the drop-down list **Key column**, select the column from the Custom DB containing the objects names as they should be saved in SOLIDserver database (string of characters: _a-z0-9 only). To prevent GUI conflicts, avoid names that are already used in the code such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*
14. In the drop-down list **Label column**, select the column from the Custom DB containing the values as they should be displayed in the list.
15. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Comment

The class object *Comment* allows to display an information, notice or warning message in the wizard among the class objects of a class. The *Comment* is always placed after the wizard's standard fields.

To add a comment

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on **Comment**. The wizard **Comment** wizard.
6. In the field **Comment**, type in the message you want to display in the wizard.
7. In the drop-down list **Style**, select the information type of comment. It can either be the content of the Comment field in a gray area (*None*), a *Notice* or a *Warning*.

Value	Description
None	The message is a comment displayed in a gray area.
Notice	The message is informational and displayed in a blue area entitled INFO.
Warning	The message is a warning displayed in an orange area entitled WARNING.

8. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

9. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. The selected comment style is displayed in the panel.

You can close Class Editor or keep adding other class objects to the same class.

Text area

The class object *Text area* allows to display a large input area that users can fill with up to 3900 characters.

To add a text area

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on **Text area**. The wizard **Text area** wizard.
6. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

7. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the field **Rows**, you can specify the number of rows to display in the text area.
 - b. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
 - c. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- d. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- e. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

10. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Horizontal separator

The class object *Horizontal separator* allows to display a line to structure and divide the class objects in the wizard.

In the creation panel, the *Horizontal separator* is represented as a straight line.

To add a horizontal separator

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on **Horizontal separator**. The wizard **Horizontal separator** opens.
6. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

7. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Jump to page

The class object *Jump to page* allows to display class objects on the new page of the wizard, you can name this page and include a message above the class objects if you want.

In the creation panel, the *Jump to page* is represented as a dotted line. Any class object located after the *Jump to page* is moved to a new page, users have to click on **NEXT** to display them.

To add a page to a wizard

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects list, click on **Jump to page**. The wizard **Jump to page** opens.
6. In the field **Title**, you can name the page of the wizard you are adding.
7. In the field **Comment**, you can type in a message that is displayed in a gray area under the page title.

8. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

You can tick the box **Translate the label** if you want the *Title* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).

9. Click on **NEXT**. The last page of the wizard opens.
10. In the drop-down list **Image**, you can select a predefined image to place on the new page.
11. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

12. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Hide IP alias

The class object *Hide IP alias* can be set on a container to skip the alias request dedicated page when assigning an IP address. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).

This object can also be set as the [Pre-defined variable](#), it corresponds to *HIDE_IP_ALIAS*.

To hide the IP alias dedicated page in the IP address addition wizard

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *IP address management*. The list refreshes.
6. In the class objects list, click on **Hide IP alias**. The wizard **Hide IP alias** opens.
7. In the field **Name**, *HIDE_IP_ALIAS* is displayed. You cannot edit it.
8. In the field **Value**, *true* is displayed. The class object is enabled.
9. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

10. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Force prefix

The class object *Force prefix* allows to force a specific prefix on a network and can be applied on the subnet-type network itself or on the block-type network or space it belongs to.

Note that there are some **special cases**:

- Forcing a prefix on a preexisting subnet-type network may cause an error.
- Forcing a prefix on a non-terminal subnet-type network only applies to the network itself, it does not apply to the terminal network it contains.

This object can also be set as the [Pre-defined variable](#), it corresponds to *FORCE_SUBNET_PREFIX*.

To force a prefix on a network

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *IP address management*. The list refreshes.
6. In the class objects list, click on **Force prefix**. The wizard **Force prefix** opens.
7. In the field **Name**, *FORCE_SUBNET_PREFIX* is displayed. You cannot edit it.
8. In the drop-down list **Network level**, select a value between 0 (block-type networks) and 15 (subnet-type networks). This value sets the level in a networks organization where the prefix is forced.
9. In the field **Value**, type in the prefix you want to force for the resource. By default, it is set to 24.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Select DHCP server

The class object *Select DHCP server* allows to display a drop-down list containing all the DHCP servers in the wizard.

To add a drop-down list containing all the DHCP servers

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DHCP management*. The list refreshes.
6. In the class objects list, click on Select DHCP server. The wizard **Select DHCP server** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows

to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select DHCP scope

The class object *Select DHCP scope* allows to display a drop-down list containing all the DHCP scopes in the wizard.

To add a drop-down list containing all the DHCP scopes

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DHCP management*. The list refreshes.
6. In the class objects list, click on **Select DHCP scope**. The wizard **Select DHCP scope** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select DHCP range

The class object *Select DHCP range* allows to display a drop-down list containing all the DHCP ranges in the wizard.

To add a drop-down list containing all the DHCP ranges

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DHCP management*. The list refreshes.
6. In the class objects list, click on . The wizard **Select DHCP range** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select DHCP static

The class object *Select DHCP static* allows to display a drop-down list containing all the DHCP statics in the wizard.

To add a drop-down list containing all the DHCP statics

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DHCP management*. The list refreshes.
6. In the class objects list, click on . The wizard **Select DHCP static** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

DHCP options

The class object *DHCP options* allows to include the DHCP options configuration fields in the wizard of a resource at the server, group, scope, range and statics level. It allows to configure DHCP options directly upon addition and edition rather than from their properties page. For more details regarding the options, refer to the appendix [DHCP Options](#).

To embed additional DHCP options

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DHCP management*. The list refreshes.
6. In the class objects list, click on DHCP options. The wizard **DHCP options** opens.
7. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a *Jump to page*.

You can set multiple conditions but they must be separated by boolean connectors.

8. Click on OK to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Select DNS server

The class object *Select DNS server* allows to display a drop-down list containing all the DNS servers in the wizard.

To add a drop-down list containing all the DNS servers

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DNS management*. The list refreshes.
6. In the class objects list, click on Select DNS server. The wizard **Select DNS server** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.

10. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:

- a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select DNS zone

The class object *Select DNS zone* allows to display a drop-down list containing all the DNS zones in the wizard.

To add a drop-down list containing all the DNS zones

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DNS management*. The list refreshes.

6. In the class objects list, click on [Select DNS zone](#). The wizard **Select DNS zone** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

12. Click on [OK](#) to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select DNS domain

The class object *Select DNS domain* allows to display a drop-down list containing all the domains, or DNS Master Name zones, in the wizard.

To add a drop-down list containing all the DNS domains

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *DNS management*. The list refreshes.
6. In the class objects list, click on **Select DNS zone**. The wizard **Select DNS zone** opens.
7. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the field **Name**, *domain* is displayed by default. You can set a different name used in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*
 - b. In the field **Label**, *Domain name* is displayed. You can set a different name for the object that will be displayed in the GUI. Only this name is seen by the user in the wizards.
 - c. Set the value of the field **Order by**, type in a value to filter the selected domain by a key. This key must respect the format: *dz.{your_value}*.
 - d. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
 - e. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
 - f. You can tick the box **Reload on change** if you want to reload the wizard page when a value is selected in the drop-down list.
 - g. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- h. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

8. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Select class

The class object *Select class* allows to include a drop-down list containing existing classes in the wizard. Only classes configured to some resources can be retrieved in this list.

To add a select class drop-down list

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Select class** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. In the drop-down list **Type**, select the resource to which apply the classes you want to display: *Space, Network, Pool, Address, DNS server, DNS zone, DHCP server, DHCP scope, DHCP range, DHCP static, User* or *Group*.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
 - b. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
 - c. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- d. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- e. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Autocompletion

The class object *Autocompletion* allows to display an autocompletion input field in the wizard. In other words, a field that provides a drop-down list returning only values that match the characters that users type in.

Under the field, a button allows to retrieve the best matches. The available values in the drop-down list depend on your configuration and can be based on from a service list or a custom DB. For more details, refer to the chapter [Custom DB](#).

- SOLIDserver services. All services and parameters are described in the API reference guide on our website ².
- A custom database.

To add an autocompletion input field using services

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.

²At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on [Autocompletion](#). The wizard **Autocompletion** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Select type**, select *Manual*.
11. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

12. Click on [NEXT](#). The next page of the wizard opens.
13. In the field **Service name**, type in the name of the listing service to call, for example *ip_subnet_list*.
14. In the field **Parameter name**, type in the name of the input parameter that should be used to pass the searched value. By default, it is *WHERE*.
15. In the field **Search condition**, type in a search condition, i.e. a variable, to display in the Autocompletion drop-down list followed by *like '%#%'*. In our example, you can type in *subnet_name like '%#%'* to format the display of all the IPv4 subnet-type networks name. You can also filter the list by replacing the hash symbol (#) by a specific matching value.
16. In the field **Parameter name for reverse search**, type in the input parameter name, used to do reverse searches. Indeed, if a user chose a network name for instance, the system only has its ID. With this parameter you can pass the ID of the object instead of a string-like parameter. By default, the parameter name is *WHERE*.
17. In the field **Reverse search condition**, type in a second variable to associate with the one to display in the drop-down list, in our example *subnet_id=#'*. You can also filter the list by replacing the hash symbol (#) by a specific matching value.
18. In the field **Key**, type in the key of the second variable, in our example *subnet_id*.
19. In the field **Display format**, type in the value that corresponds to the final display of the data in the autocompletion drop-down list. You can format this value with as many variables (preceded by \$) or literal symbols as needed. For instance, the *\$subnet_name (in \$block_name > \$site_name) - id = \$subnet_id* value displays the selected networks in the following format: *subnet_name (in block_name > site_name) - id = subnet_id*.
20. Tick the box **Allow non-matching values** if you want to allow the input field to accept values that are not part of the database.
21. Tick the box **Automatic accept** if you want the field to provide a list of matching Custom DB entries when the user types in values.

22. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add an autocompletion input field using a custom DB

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Autocompletion**. The wizard **Autocompletion** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Select type**, select *Custom DB*.
11. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

12. Click on **NEXT**. The next page of the wizard opens.
13. In the field **Custom DB name**, type in the name of the Custom DB from which you want to retrieve the data to display. The wizard refreshes.
14. In the drop-down list **Key column**, select the column containing the values to display.
15. In the drop-down list **Label column**, select the column containing the labels associated to the values to display.
16. Tick the box **Allow non-matching values** if you want to allow the input field to accept values that are not part of the database.
17. Tick the box **Automatic accept** if you want the field to provide a list of matching Custom DB entries when the user types in values.
18. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Checkbox

The class object *Checkbox* allows to insert a box in a wizard. This box is configured thanks to two values: true, when the box is ticked and false, when it is left unticked. You can use checkboxes alone or in combination with other class objects and parameters to validate complex regular expressions.

To add a checkbox

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Checkbox**. The wizard **Checkbox** opens.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. In the field **"TRUE" value**, type in the value you want to set for your box when it is ticked (value *yes* or *1*).
10. In the field **"FALSE" value**, type in the value you want to set for your box when it is not ticked (value *no* or *0*).
11. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

- a. You can tick the box **Translate the label** if you want the *Label* to be translated when users set the GUI in another language. If you tick the box, the label translation must be available on the page *Language editor*. For more details refer to the section [Customizing the Interface Names and Fields](#).
- b. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- c. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- d. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

12. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

DHCP shared network

The class object *DHCP shared network* allows to display a drop-down list containing all the scopes that can be used as shared networks in the wizard.

To associate a shared network with a resource

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **DHCP shared network** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like $\$object_value > 0$ or $\$city = "Washington"$, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Counter

The class object *Counter* allows to display a read-only field in the wizard which value increments to count the number of times the class was applied to a resource.

To add a counter

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Counter** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. You can tick the box **Padding** to display all the digits of the counter, including the zeros.
10. In the field **Number of digits**, you can type in the number of digits for your counter.
11. In the field **Min value**, you can type in the counter start value. It appears when the page is accessed for the first time.
12. In the field **Max value**, you can type in the maximum value you want to set for your counter.
13. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Force class

This class object *Force class* allows to force existing classes on IPAM resources of a lower level. If configured in a class applying to spaces, it can for instance, force classes on the networks, pools or IP addresses it contains. Note that:

- The classes you force on lower level resources must be configured and enabled
- You can force several classes on the same resource, so make sure they do not have conflicting object names.

To force a class on a lower level

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. Filter the list.
5. Make sure the class you want to force on a lower level was not set at the lowest level of the module hierarchy.
6. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
7. In the class objects drop-down list, select *All Objects*. The list refreshes.
8. In the class objects list, click on . The wizard **Force class** wizard.
9. In the drop-down list **Type**, select one of the lower levels of objects displayed according to your needs. The wizard refreshes.

If you are forcing a class on IPAM objects, selecting *Networks* displays the drop-down list *Network level*. In that list you can select a level between *0* (block-type networks) and *15* (subnet-type networks). This value sets the level in a networks organization where the class is forced.

10. In the list **Class**, double-click on the class you want to force. The class is moved to the list **Classes**.

11. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
 - b. Click on **NEXT**. The last page of the wizard opens.
 - c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.
12. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box **Required**.

You can close Class Editor or keep adding other class objects to the same class.

Force VLSM

The class object *Force VLSM* is used to force a value for the box *Terminal network* in the subnet-type networks addition wizard. When applied on a space or a block-type network, this value is set by default for all the networks they might contain. In this case, the box no longer appears in the subnet-type network addition wizard. Forcing a subnet-type network to be non-terminal enables the VLSM since it allows to create it to contain other subnet-type networks. For more details, refer to the chapter [Using VLSM to Manage Your IPAM Network](#).

This object can also be set as the [Pre-defined variable](#), it corresponds to `NO_VLSM_SUBNET`.

To force VLSM on a subnet-type network

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Force VLSM**. The wizard **Force VLSM** wizard.
7. Tick the box **Force non-terminal networks creation** if you want the class to force terminal networks to be non-terminal networks upon creation and edition. In other words, the class automatically unticks the box *Terminal network* when adding/editing subnet-type networks. If set at space level or on a block-type network, the box is not displayed or left unticked.
8. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
9. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. Select in the list **Space**, one of your an existing spaces. In a space-based VLSM organization where two sub-spaces are at the same level, this allows to favor the space you select and create the delegation within that space.

- b. Click on . The space is moved to the list **Spaces**. Only the spaces listed are available in the list *VLSM space* at the end of the wizard creating a block-type network in a child space, when a terminal network is created in a parent space.
- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

10. Click on  to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Hidden data

The class object *Hidden data* allows to associate a resource with customizable data, not displayed in the wizard. It allows to set a class object that does not require any configuration from the user or to hide another class object of the class.

This non-displayed data string could be used as a hidden signature for a class or to populate other fields if you configure the field *Constructor*, via the Expert mode.

You can also configure the class object to force a value and overwrite a preexisting content, for instance if its value is inherited from a parent. Note that all class object values, inherited or overwritten, are visible on the properties pages of the resource configured with the class.

To add hidden data

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Hidden data** wizard.
7. In the field **Name**, either type in the name of the class object in SOLIDserver database or type in the *Name* of another class object of the class.

If you specify another class object, once the Hidden data is fully configured the class object should no longer be displayed in the wizard.

8. In the drop-down list **Expert mode**, you can select **Yes** to further configure the class object:
 - a. In the field **Constructor**, you can type in the value of other class objects of the class and use them as variables to automatically overwrite the *Name* of the field with the data of your choice in the wizard.

For example, you could type in `%v{<value1>}, %v{<value2>}` where `<value#>` is the value of an existing class object in the class. If `<value1>` is a *city* and `<value2>` is a *state*, the field *Name* would be replaced with *Chicago, Illinois* in the wizard.

- b. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

- c. You can tick the box **Force the object value** to force the value of your choice when you apply the class. The field *Value to force* appears.

In the field **Value to force**, specify a value or leave it empty. The value you set is applied without being displayed in the wizard. If you specified the *Name* of another class object in the field **Name**, it overwrites the value of that object.

If you do not tick the box, the hidden data you create is configured for the class but has no value to apply.

9. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Icon

The class object *Icon* allows to display an image next to the name of a class applied to Device Manager devices, in the column **Class**.

Before configuring an *Icon* for a device class, you must upload it to the page *Custom images*. For more details, refer to the section [Uploading an Image to SOLIDserver](#).

To associate an icon with a device class

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. Filter the columns **Module** and **Type** to display *Device Manager* and *Device*.
5. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
6. In the class objects drop-down list, select *All Objects*. The list refreshes.
7. In the class objects list, click on . The wizard **Icon** opens.
8. In the field **Icon path**, type in the complete path of the icon: `/img/customImg/<uploaded-image.extension>`.
9. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

10. To display the icon next to the class name, in the sidebar, go to  **Device Manager > Devices**. The page **All devices** opens. Display the column *Device class name*. For more details refer to the section [Customizing the List Layout](#).

You can close Class Editor or keep adding other class objects to the same class.

Include class

The class object *Include class* allows to embed another class and the objects it contains. That way, a class *X* including a class *Y* that already includes a class *Z*, encompasses the objects of all three classes.

To include a class

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on [Include class](#). The wizard **Include class** opens.
7. In the drop-down list **Module**, select the module associated to the class you want to include.
8. In the drop-down list **Type**, select the type of resources associated to the class you want to include.
9. In the drop-down list **Class name**, select the class you want to include.
10. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

11. Click on [OK](#) to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Multiple input

The class object *Multiple input* allows to associate an Input with a list that displays and stores the values of the specified in input. All the values specified in the input field can be added to the list thanks to the button [ADD](#) located next to the input field.

The *Multiple input* must always be placed right under the *Input* it is associated with in the class creation panel.

Note that the *Multiple input* relies on the *Input* to be configured, but they are two different class objects. For more details, refer to the section [Select](#).

To add a multiple input field

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. Make sure an *Input* is available, without one you cannot configure a *Multiple Input*. For more details, refer to the section [Input](#).
7. In the class objects list, click on **Multiple input**. The wizard **Multiple input** wizard.
8. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

9. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
10. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
11. In the field **Input object name**, type in the name of the *Input* class object used to populate the multiple input list.
12. If you set the **Expert mode** to Yes, you can:
 - a. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- b. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- c. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows

to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

13. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.
14. Make sure the multiple input is located right under the input field. To reorder the list, refer to the section [Organizing Class Objects](#).

Once the *Multiple input* is configured, the wizard displays a button **ADD** between the input field specified in the *Input object name* and the multiple input list itself. This allows to add the value you specify in the field to the list.

You can close Class Editor or keep adding other class objects to the same class.

Multiple select

The class object *Multiple select* allows to display two lists in the wizard to select and store multiple values at the same time. The first list provides the button , to move one by one the selected values in the second list. The data listed can be:

- *Fixed values* added to the list directly from the wizard.
- *CSV values* retrieved from a CSV file.
- *Service list values* picked from the SOLIDserver services. All services and parameters are described in the API reference guide on our website ³. Note that this type of data can help you display custom databases in the wizard.

Note that a *Multiple select* is different from a *Select*. For more details, refer to the section [Select](#).

To add a multiple select drop-down list using fixed values

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Multiple select**. The wizard **Multiple select** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the wizards. This name is used for both lists. Only this name is seen by the user.

³At <http://www.efficientip.com/services/support/>, in the section *Products & Documentation*. Log in using your credentials. If you do not have credentials yet, request them at www.efficientip.com/support-access.

9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Select type**, select *Fixed values*. The wizard refreshes.
11. In the field **Key**, type in an object name for the database using only the string of characters *_a-z0-9*⁴. The label **Label/Key** displays the *Key*.
12. In the field **Label**, type in the label of the the *key* that is displayed in the list. The field **Label/Key** autopopulates as follows: *<Key>#<Label>*.
13. Click on **+**. The value is moved to the list **Options**.
14. Repeat these actions for as many values as needed. You can use **-** to remove one by one values from the list, or **↕** and **↔** to reorganize them.
15. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
 - b. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.
 - c. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- d. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- e. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

⁴To prevent GUI conflicts, avoid names that are already used in SOLIDserver database like *site*, *mac-addr*, *gateway*, *vlan*, *domain*, *user*, *port*, *password*...

16. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add a multiple select drop-down list using CSV values

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Multiple select** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the wizards. This name is used for both lists. Only this name is seen by the user.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Select type**, select *CSV values*. The wizard refreshes.
11. In the field **CSV file**, type in the complete path of the file stored in the appliance.
12. In the field **Value column**, type in the number of the column in the CSV file containing the values to retrieve.
13. In the field **Label column**, type in the number of the column in the CSV file containing the labels corresponding to the values to retrieve.
14. In the field **Filter column**, type in the number of the column used to match certain rows.
15. Click on . The value is moved to the list **Filter**. You can use to remove the values from the list one by one, or and to reorganize them.
16. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:
 - a. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
 - b. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.
 - c. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- d. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- e. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

17. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

To add a multiple select drop-down list using services

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Multiple select** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the wizards. This name is used for both lists. Only this name is seen by the user.
9. You can tick the box **Required** to make the class object configuration compulsory in the resource addition and edition wizards.
10. In the drop-down list **Select type**, select *Service list values*. The wizard refreshes.
11. In the field **Services**, start typing in the name of the service to call. The matching services are listed, select the one that suits your needs.

12. In the field **Key**, type in the object name as it should be saved in SOLIDserver database (string of characters: `_a-z0-9` only). To prevent GUI conflicts, avoid names that are already used in the code such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*
13. In the field **Label**, type in the name of the input parameter corresponding to the labels associated to these values.
14. In the field **Where**, type in an SQL condition to filter the retrieved values if need be.
15. In the field **Order by**, type in an SQL condition to sort the results if need be.
16. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

- a. You can tick the box **Have none label** if you want to include the value *None* to the drop-down list.
- b. In the field **Default value**, you can type in a default value for the field if no data is specified. The value you specify is by default displayed in the wizard and users can edit it.
- c. In the drop-down list **Inheritance property**, you can configure the inheritance behavior of the value of the class object:

Value	Description
None	The property is not set, this is the default value. Users can set it to <i>Inherit</i> or <i>Set</i> .
Inherit	The property is forced to <i>Inherit</i> . Users cannot change it to <i>None</i> or <i>Set</i> .
Set	The property is forced to <i>Set</i> . Users cannot change it to <i>None</i> or <i>Inherit</i> .

For more details regarding the inheritance property, refer to the chapter [Inheritance and Propagation](#).

- d. In the drop-down list **Propagation property**, you can configure the propagation behavior of the value of the class object:

Value	Description
None	The property is not set, this is selected by default. Users can set it to <i>Propagate</i> or <i>Restrict</i> .
Propagate	The property is forced to <i>Propagate</i> . Users cannot change it to <i>None</i> or <i>Restrict</i> .
Restrict	The property is forced to <i>Restrict</i> . Users cannot change it to <i>None</i> or <i>Propagate</i> .

For more details regarding the propagation property, refer to the chapter [Inheritance and Propagation](#).

- e. In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

17. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel. It is followed by  if you ticked the box *Required*.

You can close Class Editor or keep adding other class objects to the same class.

Object name

The class object *Object name* allows to build an automatic naming rule for a resource, such as `%v{city}-%v{store code}` where *city* and *store code* are the names of objects belonging to the same class. By convention, an *Objectname* and the class objects used to build it should be placed in the first page of the wizard.

To name a resource automatically

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Objectname** opens.
7. In the field **Constructor**, you can type in the value of other class objects of the class and use them as variables to automatically overwrite the *Name* of the field with the data of your choice in the wizard.

For example, you could type in `%v{<value1>}, %v{<value2>}` where `<value#>` is the value of an existing class object in the class. If `<value1>` is a *city* and `<value2>` is a *state*, the field *Name* would be replaced with *Chicago, Illinois* in the wizard.

8. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

9. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Pre-defined variable

The class object *Pre-defined variable* allows to force some behaviors and configuration choices in the wizard. The pre-defined variables available can be applied to the user rights and to the modules: IPAM, DHCP, DNS, Device Manager and Workflow.

Each pre-defined variable must be configured via the field *Value* that can accept specific values or 1 to be activated. To understand the purpose of each variable, refer to the appendix [Class Studio Pre-defined Variables](#).

To insert a pre-defined variable

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Pre-defined variable**. The wizard **Pre-defined variable** opens.
7. Configure the variable, following the details in the appendix [Class Studio Predefined Variables](#).
 - a. In the drop-down list **Name**, select the predefined variable of your choice.
 - b. In the field **Value**, type in the value that suits your needs.
8. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.
9. Click on **OK** to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Time stamp

The class object *Time stamp* allows to include the exact time and date of creation of a resource in a dedicated field of the wizard.

Note that if you edit a resource that was not configured with this class object, the time stamp matches the moment you apply the class and not its actual creation date and time.

To add a field Time stamp

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on **Time stamp**. The wizard **Time stamp** wizard.
7. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like `$object_value > 0` or `$city=="Washington"`, that allows to only display the class object if your condition is matched. Note that the condition can only be

taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

8. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Upload file

The class object *Upload file* allows to include an upload tool in the wizard. It displays a field **File name** and a button that opens a window and allows to upload any file to SOLIDserver database. Keep in mind that:

- Uploaded files cannot exceed 300 MB.
- Uploaded files are stored temporarily in the */tmp* folder of the appliance and deleted shortly after. The upload tool can therefore be used to import CSV files or other types of files to be processed straight away by other class objects.

To add a field Upload file

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Upload file** wizard.
7. In the field **Name**, type in the name of the class object in SOLIDserver database.

You can use hexadecimal characters and underscores "_", but no spaces. To prevent GUI conflicts, avoid names that are already used within the appliance such as: *site, mac-addr, gateway, vlan, domain, user, port, password...*

8. In the field **Label**, type in the class object name displayed in the GUI. Only this name is seen by the user in the wizards.
9. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

10. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Owner

The class object *Owner* allows to display the name of the user who added the resources, no matter who edits it later.

To add a field Owner

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the class objects drop-down list, select *All Objects*. The list refreshes.
6. In the class objects list, click on . The wizard **Owner** wizard.
7. In the drop-down list **Expert mode**, you can select *Yes* to further configure the class object:

In the field **Show if...**, you can condition the display of the class object in the wizard in the form of an "if" statement, like *\$object_value > 0* or *\$city=="Washington"*, that allows to only display the class object if your condition is matched. Note that the condition can only be taken into account if the value of the class object has already been saved in the wizard, either via inheritance or because it is located after a [Jump to page](#).

You can set multiple conditions but they must be separated by boolean connectors.

8. Click on to complete the operation. The object is now embedded into the class and listed in the creation panel.

You can close Class Editor or keep adding other class objects to the same class.

Editing Class Objects

Class objects can be edited at any time even if the class they belong to is already in use.

Keep in mind that **renaming an object already used by a resource will delete all the class data it is associated with**. It can only be retrieved by renaming the object back, before filling any new class data through the newly edited object.

To edit a class object

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on ✨ **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the creation panel, click on the gray box displaying the name of the class object to edit. The corresponding object class wizard opens.
6. Edit the class object according to your needs following the procedure that suits your needs in the section [Adding Class Objects](#).

7. Click on to complete the operation. The object is updated in the creation panel.

Organizing Class Objects

The way user defined fields are displayed on the Add/Edit wizards of a resource depends on the organization of the class objects when configuring the class. SOLIDserver simplifies this process by allowing members of the *admin* group to drag & drop the objects of their choice.

Keep in mind that a *Multiple select* can only be effective if placed underneath an *Input*.

To organize class objects

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the creation panel, drag and drop the class objects to change their display order once the class is used on a resource.

Deleting Class Objects

At any time, you can remove a class object from a class. Keep in mind that deleting a class object removes the related user defined fields and data form all the resources wizards it is applied to.

To delete a class object

Only users of the group *admin* can perform this operation.

1. Make sure your browser allows pop-up windows.
2. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
3. In the section **Customization**, click on **Class Studio**. The page **Class Studio** opens.
4. In the column **Name**, click on the class you want to edit. **Class Editor** opens.
5. In the creation panel, click on  next to the class object you want to delete. The wizard opens.
6. Click on to complete the operation. The object is no longer listed in the creation panel.

Class Studio Syntax

Class Studio is a powerful tool to help members of the *admin* group in provisioning SOLIDserver through classes that act as automated rules. Configuring these classes requires knowledge of the regular expression (regex) syntax. Regex being as powerful as complex, this section presents only the basic information to help you configure simple class parameters. For more details regarding regular expressions, you can visit the websites <http://www.regular-expressions.info> and <http://regexlib.com>.

Basic Regular Expressions

regex stands for "regular expression". Regular expressions are used by the administrator to enforce the format of the data input in the user defined fields. For instance, it can be used to set a maximum string length, allow only alpha characters a-z or A-Z or numbers in the range [1-99].

For instance, in the following example of regex: **^[a-zA-Z]{1,10}\$**

^ means "begin matching at start of string"

[a-zA-Z] means "match lower case and upper case letters a-z"

{1,25} means "match the previous item (the letters which format was explained in the previous line) 1 to 25 times"

\$ means "only match if cursor is at end of string"

Basic regular expression symbols are described in the table below:

Table 97.6. Basic regex symbols

Symbol	Description
.	Matches any single character. If put between bracket, the dot symbol matches a literal dot. For example, a.c matches "abc", etc., but [a.c] matches only "a", ".", or "c".
[]	A bracket expression. Matches a single symbol contained within the brackets. For example, [abc] matches "a", "b", or "c". [a-z] specifies a range matching any lowercase letter from "a" to "z". These forms can be mixed: [abcx-z] matches "a", "b", "c", "x", "y", or "z", as does [a-cx-z]. The - symbol is treated as a literal symbol if it is the last or the first (after the ^) symbol within the brackets: [abc-], [-abc]. Note that backslash escapes are not allowed. The] symbol can be included in a bracket expression if it is the first (after the ^) symbol: []abc].
[^]	Matches a single symbol that is not contained within the brackets. For example, [^abc] matches any symbol other than "a", "b", or "c". [^a-z] matches any single symbol that is not a lowercase letter from "a" to "z". Likewise, literal symbols and ranges can be mixed.
^	Matches the starting position within the string. In line-based tools, it matches the starting position of any line.
\$	Matches the ending position of the string or the position just before a string-ending newline. In line-based tools, it matches the ending position of any line.
()	Defines a marked subexpression.
*	Matches the preceding element zero or more times. For example, ab*c matches "ac", "abc", "abbbc", etc. [xyz]* matches "", "x", "y", "z", "zx", "zyx", "xyzzy", and so on. (ab)* matches "", "ab", "abab", "ababab", and so on.

Defining a Class as a Group Resource

In SOLIDserver, only users of the group *admin* can manage and modify the items of every module. Adding a customized class as one of the resources of a specific group allows the users of that group to apply the class as long as they have the corresponding rights granted.

Granting access to a class as a resource also grants access to its class parameters, so users can configure them when they add or edit the objects configured with the class. If a group does not have a customized class in its resources, editing objects configured with it empties the value of all the class parameters. For more details, refer to the section [Assigning Resources to a Group](#) in the chapter Managing Groups.

Chapter 98. Configuring Custom Databases

You can create custom databases from the page **Custom DB**. These databases are directly embedded in SOLIDserver and can contain a maximum of 10 pieces of information, columns, named *Label #*.

By default, a custom DB named **Vendor** is already installed. This database is used by SOLIDserver in order to link MAC address and the Vendor of the Ethernet card together. It cannot be modified. This is an example of what a custom DB could be.

Keep in mind that the custom databases can come in handy when configuring classes with parameters like select, multiple select or autocompletion for instance. For more details, refer to the chapter [Configuring Classes](#).

Browsing Custom DB

Custom DB is divided into two pages: one displaying the databases themselves and the other displaying the data of each custom database.

Browsing Custom DB Database

To display the list of custom databases

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The **Custom database** listing page opens. By default, the custom DB database contains the database *Vendor*.

To display a custom database properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. At the end of the line of the custom database of your choice, click on . The properties page opens.

On the properties page, the panel **Main properties** displays the Custom database name, Type, Description and labels it contains.

Browsing Custom Data

To display the list of custom data

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the breadcrumb, click on **Custom data**. The page **Custom data** opens.

To display the custom data of a specific custom DB

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the column **Name**, click on the name of the custom database of your choice. The page **Custom data** opens, it only displays the data of the custom database you chose.

To display a custom data entry properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. At the end of the line of the entry of the custom data of your choice, click on . The properties page opens.

On the properties page, the panel **Main properties** displays the name of the Custom database it contains to and its defined labels and their value.

Adding a Custom DB

To add a custom DB

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the menu, click on  **Add**. The wizard **Create a custom DB** opens.
4. Fill in each field as describe below:

Table 98.1. CustomDB creation parameters

Field	Description
Database name	Choose a name for the custom DB you are creating. This field is required.
Type	Define the database type.
Description	Describe the database. This field is optional.
Label 1 to Label 10	Name the columns of your custom database.

5. Click on  to complete the operation.

Editing a Custom DB

At any time you can edit a custom database. However, keep in mind that:

- You should not rename a database with an already used name.
- The default custom database *Vendor* is the only one in Read-only. It cannot be edited.

To edit a custom DB

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the column **Name**, put your mouse over the name of the database you want to edit. Click on . The wizard **Edit custom database** opens.
4. Edit each field according to your needs following the table below:

Table 98.2. CustomDB parameters

Field	Description
Name	Choose a name for the custom DB you are editing. This field is required.
Type	Define the database type.
Description	Describe the database. This field is optional.
Label 1 to Label 10	Name the columns of your custom database.

- Click on to complete the operation. The wizard refreshes and closes. The changes are displayed in the list.

Deleting a Custom DB

To delete a custom DB

- In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
- In the list, tick the custom database that you want to delete.
- In the menu, click on  **Delete**. The wizard opens.
- Click on to complete the operation.

Configuring a Custom DB with Custom Data

Once a custom database is created, it is empty and you must add custom data.

Once your database contains everything you need, you can use it within classes and apply it to the resource that suits your needs. For more details, refer to the chapter [Configuring Classes](#).

Adding Data in a Custom DB

From the page *Custom data*, you can add entries to existing custom databases.

You can also import custom data from a CSV file. For more details, refer to the section [Importing Custom Data](#).

To add a custom data entry from the page Custom data

- In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
- In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
- In the breadcrumb, click on **Custom data**. The page **Custom data** opens.
- In the menu, click on  **Add**. The wizard **Add custom data** opens.
- In the drop-down list **Database**, select the custom database of your choice.
- Click on . The last page of the wizard opens.
- Fill in the fields according to your needs, all 10 fields are optional. The fields can either have the default name *Value 1... Value 10* or have the name (label) that was set for custom database.
- Click on to complete the operation.

To add a custom data entry in a specific custom DB

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the list, click on the name of the custom database of your choice. The page **Custom data** opens.
4. In the menu, click on **+ Add**. The wizard **Add custom data** opens.
5. Fill in the fields according to your needs, all 10 fields are optional. The fields can either have the default name *Value 1... Value 10* or have the name (label) that was set for custom database.
6. Click on to complete the operation.

Editing the Data of a Custom DB

You can edit the content of a custom DB. However, like for the custom DB, you should not rename data with a name already used.

To edit custom data

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the column **Name**, click on the custom database of your choice. The page **Custom database** opens.
4. In the column **Name**, put your mouse over the data you want to edit. Click on . The wizard opens.
5. Edit the value of the fields according to your needs.
6. Click on to complete the operation. The wizard refreshes and closes. The changes are displayed in the list.

Deleting Data From a Custom DB

To remove data from a custom DB

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Custom DB**. The page **Custom database** opens.
3. In the column **Name**, click on the custom database of your choice. The page **Custom database** opens.
4. Tick the custom data entries that you want to remove.
5. In the menu, click on  **Delete**. The wizard opens.
6. Click on to complete the operation.

Chapter 99. Managing Customization Packages

From the Administration module, you can import archive files, or packages, containing a set of customized functionalities from the page **Packager**.

Once imported and uploaded, installing these packages can affect interfaces, databases, system files, etc. depending on what they contain. These functionalities can take the form of classes, services (also called macros), reports or rules.

Packager is composed of two pages: All Packages and All package files. From the page All Packages, you can import or create, install, uninstall and delete your packages. The page All package files simply provides the content of the packages.

Packager reuses the principle of the module of the same name in 3.0.1, however, it uses different services. Therefore, packages created or used in previous versions of SOLIDserver cannot be used with the current version.

Browsing the Packages Database

The packages and their content are displayed on two different pages. The packages management options are only available on the page *All packages*.

To display the list of packages

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.

The page *All packages* contains seven columns: **Name**, **Description**, **Version**, **Vendor**, **Creation time**, **Install time** and **Status**. You cannot edit the page listing template.

To display all the information in one panel, open the package properties page.

To display a package properties page

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. At the end of the line of the package of your choice, click on . The server properties pages opens.

The packages content is listed on the page *All package files*.

To display the list of package files

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. In the breadcrumb, click on **All package files**. The page **All package files** opens.

To display the package files of a specific package

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. Click on the name of the package of your choice. The page **All package files** of the selected package opens.

The page *All package files* contains five columns: **filename**, **Directory**, **Type**, **Package version** and **Version**. You cannot edit the page listing template.

Uploading Packages

From the page All packages, you can upload your own packages in a .tar archive file.

Keep in mind that:

- Uploading a package simply stores it locally on the appliance. Once uploaded, you need to install it to push the files it contains. For more details, refer to the section [Installing Packages](#).
- Each package has a unique name, version and content, so you cannot upload a package if it is already listed on the page unless the version or name differs. If at least one of the files it contains is already installed, you cannot install your package.
- Packages from previous versions of SOLIDserver are not compatible and therefore not supported.

To upload a package

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. In the menu, click on **+ Add**. The wizard **Upload a package** opens.
4. Click on **BROWSE** to search for the .tar file to import. A window opens to help you browse through folders.
5. Double-click on the needed file. The window closes and the file is visible in the field **File name** of the wizard.
6. Click on **OK** to complete the operation. The report opens and closes. The page **All Packages** is visible again. The package is listed but it is not installed yet.

Creating Packages

If you want you can create your own packages from the page All packages. In this case, you can configure it with existing rules, services, reports and classes.

Keep in mind that:

- Creating a package does not install it. Once created, you need to install it to push the files it contains. For more details, refer to the section [Installing Packages](#).
- Each package has a unique name, version and content, so you cannot upload a package if it is already listed on the page unless the version or name differs. If at least one of the files it contains is already installed, you cannot install your package.

- You cannot include system files to your package. If you include any of SOLIDserver system files during the creation, you cannot install the package.

To create a package

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. In the menu, select  **Tools** > **Expert** > **Create a package**. The wizard **Create a package** opens.
4. Configure the package details:
 - a. In the field **Package name**, name the package.
 - b. In the field **Version**, type in a version for your package following the format `<number>.<number>`.
 - c. In the field **Description**, you can describe the package.
 - d. In the field **Vendor**, you can type in a vendor name.
5. Click on **NEXT** to configure the package content. The page **Package files selection** opens.
6. You can add classes to your package:
 - a. In the drop-down list **Files type**, select *Class*.
 - b. In the drop-down list **Module**, select the module of your choice. It can be *Administration, DHCP, DNS, Device Manager, IPAM, NetChange, Rights & delegation, SPX, VLAN Manager, VRF* or *Workflow*.
 - c. In the drop-down list **Type**, select the object the class applies to within the module.
 - d. In the drop-down list **Available files**, select an existing class. If it belongs to a specific directory, it is listed as follows: `<directory-name>/<class-name>`. For more details regarding classes, refer to the chapter [Configuring Classes](#).
 - e. Once you selected the class that suits your needs, click on **ADD**. The class is moved to the list **Selected files**. You can add as many classes as needed.
7. You can add services to your package:
 - a. In the drop-down list **Files type**, select *Macro*.
 - b. In the drop-down list **Module**, select the module of your choice. It can be *Administration, DHCP, DNS, Device Manager, IPAM, NetChange, Rights & delegation, SPX, VLAN Manager, VRF* or *Workflow*.
 - c. In the drop-down list **Available files**, select the service of your choice.
 - d. Once you selected the service that suits your needs, click on **ADD**. The service is moved to the list **Selected files**. You can add as many services as needed.
8. You can add reports to your package:
 - a. In the drop-down list **Files type**, select *Report*.
 - b. In the drop-down list **Available files**, select the report of your choice. For more details regarding the reports refer to the chapter [Managing Reports](#).

- c. Once you selected the report that suits your needs, click on **ADD**. The report is moved to the list **Selected files**. You can add as many reports as needed.
9. You can add rules to your package:
 - a. In the drop-down list **Files type**, select *Rule*.
 - b. In the drop-down list **Module**, select the module of your choice. It can be *Administration, DHCP, DNS, Device Manager, IPAM, NetChange, Rights & delegation, SPX, VLAN Manager, VRF* or *Workflow*.
 - c. In the drop-down list **Available files**, select an existing rule applying to the selected module.
 - d. Once you selected the rule that suits your needs, click on **ADD**. The rule is moved to the list **Selected files**. You can add as many rules as needed.
10. In the drop-down list **Selected files** are listed all the classes, services, reports and rules that the package contains.

To remove an entry, select it the list and click on **DELETE**.
11. Click on **OK** to complete the operation. The report opens and closes. The page **All Packages** is visible again. The package is listed but it is not installed yet.

Editing Packages

You **cannot edit a package**. If one of your packages contains files that you no longer require or if it misses files, **you need to replace it**.

1. Uninstall the useless package.
2. Upload the package that replaces it or create another package. To make sure you do not forget any file, you can look at the list *All package files* of the package you want to replace.
3. Delete the useless package.
4. Install the new package.

Installing Packages

Installing a package pushes its files to the relevant parts of the appliances. When uploading or creating a package, it is simply listed in the GUI. If you do not install it, the files it contains are simply stored locally but not used.

Keep in mind that:

- You cannot install a package containing SOLIDserver system files.
- You cannot install a package if it contains one or several files that were already installed with another package.
- Once you installed a package, you cannot delete it, you must uninstall it before being able to delete it.

To install a package

1. In the sidebar, click on **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. Tick the package(s) you want to install.
4. In the menu, select  **Edit** > **Install**. The wizard **Install a package** opens.
5. Click on to complete the operation. The report opens and works until all the files are pushed. The page is visible again. In the column **Status** the package is marked  *installed*.

Uninstalling Packages

Uninstalling a package allows to revert all the changes that the files it contains were performing. It also allows to delete a package: you cannot delete a package if it is installed, i.e. used.

To uninstall a package

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. Tick the package(s) you want to uninstall.
4. In the menu, select  **Edit** > **Uninstall**. The wizard **Uninstall a package** opens.
5. Click on to complete the operation. The report opens and closes. The page is visible again. In the column **Status** the package is marked  *uninstalled*.

Downloading Packages

At any time you can download a package, whether it is installed or not. Keep in mind that you can only download one package at a time.

To download a package

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. Tick the package you want to download.
4. In the menu, select  **Edit** > **Download**. The wizard **Downloading a package** opens.
5. Click on to complete the operation. The report opens and displays the package which is an archive .tar file that can be downloaded from the page **Local files listing** available from the page *Admin Home*.
6. Click on to download the file before closing the wizard.
7. Click on . The wizard closes and the page **All packages** is visible again.

Deleting Packages

Once you no longer need a package, you can delete it as long as it is no longer used. This means that if the package you want to delete is currently installed, you need to uninstall it before following the procedure below. For more details, refer to the section [Uninstalling Packages](#).

To delete a package

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.

2. In the section **Customization**, click on **Packager**. The page **All packages** opens.
3. Tick the package(s) you want to delete.
4. In the menu, click on  **Delete**. The wizard **Delete** opens.
5. Click on to complete the operation. The report opens and closes. The package is no longer listed.

Appendix A. Matrices of Network Flows

Table of Contents

SOLIDserver	1234
IPAM	1235
DHCP	1236
DNS	1237
NetChange	1239
Remote Management	1240

This appendix maps out the networks flows that you must open to manage your SOLIDserver appliance or remotely manage servers.

Each flow is detailed in tables indicating its: **Source IP, Port, Destination IP, Port, Protocol, Service** used and some **Notes**, when relevant.

The columns *Source IP* and *Destination IP* may contain the following key words:

Source/Destination	Description
administrator	The computer belonging to a user with administrative rights.
DHCP server	Any DHCP server, the flow must be opened for all servers.
DHCP backup	A server managed in a SOLIDserver smart architecture that has a backup role in your failover configuration.
DHCP master	A server managed in a SOLIDserver smart architecture that has a master role in your failover configuration.
DNS server	Any DNS server, the flow must be opened for all servers.
HSM	Your Hardware Security Module.
iDRAC	The integrated Dell Remote Access Controller for hardware appliances SDS-260 to SDS-Blast Series.
Kerberos servers	Your Kerberos authentication server.
MS DHCP	A Microsoft DHCP server managed from SOLIDserver.
MS DNS	A Microsoft DNS server managed from SOLIDserver.
Network device	All the routers, switches and/or firewalls which information you want to retrieve. They must be managed in the module NetChange.
RFS	The Remote File System of your HSM.
SOLIDserver	Any SOLIDserver appliance.
SOLIDserver Hot Standby	A SOLIDserver appliance configured in High availability which role is Hot Standby. As the Hot Standby can be switched to Master, the matrices flows on both HA appliances should be configured the same.
SOLIDserver Management	A SOLIDserver appliance managing another SOLIDserver or any external server or service.
SOLIDserver Master	A SOLIDserver appliance configured in High availability which role is Master. The matrices flows on both the Master and Hot Standby appliances should be configured the same.

SOLIDserver

Basic Configuration

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
administrator	any	SOLIDserver	80	TCP	HTTP	Graphic User Interface (GUI)
administrator	any	SOLIDserver	443	TCP	HTTPS	Graphic User Interface (GUI)
administrator	any	SOLIDserver	22	TCP	SSH	Command Line Interface (CLI)
SOLIDserver	any	DNS server	53	UDP	DNS	DNS resolution, DDNS update
SOLIDserver	any	DNS server	53	TCP	DNS	DNS resolution, DNS zone transfer
SOLIDserver	any	NTP server	123	UDP	NTP	Time synchronization
SOLIDserver	any	FTP server	21	TCP	FTP	Remote archive on an FTP or SFTP server
SOLIDserver	any	SFTP server	22	TCP	SFTP	

External Authentication

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver	any	LDAP/AD	389	TCP	LDAP	LDAP or AD authentication
SOLIDserver	any	LDAPS/AD	636	TCP	LDAPS	LDAPS or AD authentication
SOLIDserver	any	RADIUS	1812	UDP	RADIUS	RADIUS authentication

TFTP and NTP Services

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
TFTP client	any	SOLIDserver	69	UDP	TFTP	File transfer
NTP client	any	SOLIDserver	123	UDP	NTP	NTP server

Monitoring and Logging

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
Monitoring server	any	SOLIDserver	161	UDP	SNMP	SNMP pooling
SOLIDserver	any	SOLIDserver	161	UDP	SNMP	
SOLIDserver	any	Monitoring server	162	UDP	SNMP	SNMP trap
SOLIDserver	any	Log server	514	UDP	Syslog	Syslog export

iDRAC

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
administrator	any	iDRAC	22	TCP	SSH	iDRAC SSH
administrator	any	iDRAC	80	TCP	HTTP	iDRAC GUI
administrator	any	iDRAC	443	TCP	HTTPS	iDRAC GUI
administrator	any	iDRAC	5900	TCP	VNC	Virtual Console

IPAM

SPX

RIPE

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver	any	RIPE database	43	TCP	WHOIS	Retrieve RIPE data (Whois)
SOLIDserver	any	RIPE database	80	TCP	WHOIS	Send data to the RIPE (Whois)
SOLIDserver	any	SMTP server	25	TCP	SMTP	Send data to the RIPE via a mail server
SOLIDserver	any	POP server	110	TCP	POP	
SOLIDserver	any	RIPE database	443	TCP	HTTPS/POST	Send data to the RIPE directly

APNIC

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver	any	APNIC data-base	43	TCP	WHOIS	Retrieve APNIC data (Whois)
SOLIDserver	any	APNIC data-base	80	TCP	WHOIS	Send data to the APNIC (Whois)
SOLIDserver	any	SMTP server	25	TCP	SMTP	Send data to the APNIC via a mail server
SOLIDserver	any	POP server	110	TCP	POP	
SOLIDserver	any	APNIC data-base	443	TCP	HTTPS/POST	Send data to the APNIC directly

DHCP

EfficientIP DHCP Servers

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DHCP server	443	TCP	HTTPS	Required to manage an EfficientIP DHCP server on a SOLIDserver appliance
DHCP master	any	DHCP backup	647	TCP	Failover	Failover channel port on the backup server
DHCP backup	any	DHCP master	847	TCP	Failover	Failover channel port on the master server
DHCP client	68	DHCP server	67	UDP	DHCP	Required by the service DHCP
DHCP server	67	DHCP client	68	UDP	DHCP	Required by the service DHCP
DHCP client	546	DHCP server	547	UDP	DHCP	Required by the service DHCPv6
DHCP server	547	DHCP client	546	UDP	DHCP	Required by the service DHCPv6
DHCP client	68	Broadcast address	67	UDP	DHCP	Required by the DHCP protocol on the local segment
DHCP server	-	any	-	ICMP	ICMP	Only if the option <i>ping-check</i> ^a is enabled

^aFor more details, refer to the section [Preventing IP Address Duplication](#).

Microsoft Windows Agentless DHCP Servers

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	MS DHCP	135	TCP	MSRPC	Microsoft Remote Procedure Calls (MSRPC)
DHCP client	68	MS DHCP	67	UDP	DHCP	Required by the service DHCP
MS DHCP	67	DHCP client	68	UDP	DHCP	Required by the service DHCP
DHCP client	68	Broadcast address	67	UDP	DHCP	Required by the DHCP protocol on the local segment

Linux Packages

Prerequisite before configuring a Linux Package: configuring DHCP network flows as detailed in the section [EfficientIP DHCP Servers](#).

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DHCP server	443	TCP	HTTPS	Required to manage the DHCP server on Linux packages

DHCP Statistics

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DHCP server	161	UDP	SNMP	SNMP v1, v2c and v3 to retrieve the server statistics

DNS

EfficientIP DNS Servers

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DNS server	443	TCP	HTTPS	Required to manage an EfficientIP DNS server on a SOLIDserver appliance
SOLIDserver Management	any	DNS server	53	UDP/TCP	DNS	DNS resolution, DDNS update, DNS zone transfer
DNS server	any	DNS server	53	UDP/TCP	DNS	DNS resolution, DDNS update, DNS zone transfer
DNS server	53	SOLIDserver Management	2053 ^a	UDP	DNS	DNS notify (optional)

^aThe port number 2053 is used to send notify from the DNS server to the management platform. This notify can be configured to speed up the RR upload on DNS zone change. Keep in mind that not all the DNS engines support this functionality, for instance Microsoft DNS engines do not support it.

Microsoft Windows Agentless DNS Servers

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	MS DNS	135	TCP	MSRPC	Microsoft Remote Procedure Calls (MSRPC)
SOLIDserver Management	any	MS DNS	53	UDP/TCP	DNS	DNS resolution, DDNS update, DNS zone transfer
MS DNS	any	MS DNS	53	UDP/TCP	DNS	DNS resolution, DNS zone transfer

Amazon Route 53 Servers

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	AWS Route 53	443	TCP	HTTPS	Required to manage a Route 53 DNS server on Amazon Web Service (AWS)
SOLIDserver Management	any	AWS Route 53	53	UDP/TCP	DNS	DNS resolution, DDNS update, DNS zone transfer

Linux Packages

Prerequisite before configuring a Linux Package: configuring the DNS network flows as detailed in the section [EfficientIP DNS Servers](#).

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DNS server	443	TCP	HTTPS	Required to manage the DNS server on Linux packages

DNS Statistics

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	DNS server	161	UDP	SNMP	SNMP v1,v2c or v3 to retrieve the statistics of DNS servers on SOLIDserver appliances or Linux Packages

GSS-TSIG

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	Kerberos servers	88	Protocol	Kerberos	Kerberos authentication

Routing Protocols for Anycast

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
Router's IP(s)	any	SOLIDserver	179	TCP	BGP	
Router's IP(s)	-	224.0.0.0/24	-	OSPF	OSPF	
SOLIDserver	-	224.0.0.0/24	-	OSPF	OSPF	
-	-	-	-	-	IS-IS	

Guardian Cache Sharing

Via Unicast

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver	any	SOLIDserver	<i><port set by user></i>	UDP	Cache Sharing	Guardian cache sharing, via the port of your choice

Via Multicast

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
Local Network(s)	-	224.0.0.0/24	-	IGMP	IGMP	IGMP
SOLIDserver	-	224.0.0.0/24	-	IGMP	IGMP	

HSM

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
DNS server	any	HSM	9004	UDP/TCP	Thales HSM	Required for DNSSEC signing with HSM
HSM	any	RFS	9004	TCP	RFS	
DNS server	any	RFS	9004	TCP	RFS	

NetChange

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
NetChange	any	Network device	161	UDP	SNMP	SNMP v1, v2c, v3
NetChange	any	DNS server	53	UDP	DNS	DNS resolution
NetChange	any	Network device	22	TCP	SSH	Save the configuration
NetChange	any	Network device	23	TCP	SNMP	

Remote Management

High Availability

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Hot Standby	any	SOLIDserver Master	443	TCP	HTTPS	Health check
SOLIDserver Master	any	SOLIDserver Hot Standby	443	TCP	HTTPS	
SOLIDserver Hot Standby	any	SOLIDserver Master	5432	TCP	PostgreSQL	Replication
SOLIDserver Master	any	SOLIDserver Hot Standby	5432	TCP	PostgreSQL	

Remote Management of Other Appliances

Source IP	Port	Destination IP	Port	Protocol	Service	Notes
SOLIDserver Management	any	SOLIDserver	443	TCP	HTTPS	Management of remote SOLIDserver appliances

Appendix B. Multi-Status Messages

Table of Contents

DHCP Multi-Status Messages	1241
DNS Multi-Status Messages	1241

This appendix provides a list of messages returned in the column **Multi-status**. For more details, refer to the section [Understanding the Column Multi-Status](#).

DHCP Multi-Status Messages

Code	Level	Message
60000	Multistatus	60000: Communications-interrupted.
60001	Multistatus	60001: Partner-down.
60002	Multistatus	60002: Recovering.
60003	Multistatus	60003: Starting up.
60004	Multistatus	60004: Server management via SNMP can only be in read-only. We recommend to update to SSL.
60005	Multistatus	60005: This server type is no longer supported. To manage a Microsoft server, create a "Microsoft DHCP" server.

DNS Multi-Status Messages

Code	Level	Message
61000	Multistatus	61000: Zone type incompatible with Hybrid.
61001	Multistatus	61001: Hybrid servers cannot manage authoritative and recursive zones.
61002	Multistatus	61002: Hybrid does not support forward on authoritative servers.
61003	Multistatus	61003: Hybrid does not support forwarders on authoritative servers.
61004	Multistatus	61004: Hybrid server with authoritative zones cannot be recursive.
61005	Multistatus	61005: RR type incompatible with Hybrid.
61006	Multistatus	61006: Server type incompatible with Hybrid.
61007	Multistatus	61007: TSIG keys not supported on Hybrid recursive servers.
61008	Multistatus	61008: Hybrid servers do not support views.
61009	Multistatus	61009: Hybrid does not support forwarding configuration on authoritative servers.
61010	Multistatus	61010: Zone type incompatible with Route 53.
61011	Multistatus	61011: RR type incompatible with Route 53.
61012	Multistatus	61012: Route 53 servers do not support views.
61013	Multistatus	61013: At least one character in the value of the record is not supported by Route 53 servers.
61014	Multistatus	61014: Only TLD zones are replicated on Route 53 servers.
61015	Multistatus	61015: Maximum number of AWS zones reached.
61016	Multistatus	61016: Maximum number of RRset per AWS zone reached.
61017	Multistatus	61017: Maximum number of records per RRset per AWS zone reached.

Multi-Status Messages

Code	Level	Message
61018	Multistatus	61018: This RPZ zone cannot be replicated on one of the physical servers of the smart.
61019	Multistatus	61019: The syntax of the BIND include file is incorrect.
61020	Multistatus	61020: RRL is not supported on this version of BIND.
61021	Multistatus	61021: RRL is not supported on this DNS server.
61022	Multistatus	61022: The zone name does not comply with AWS format.
61023	Multistatus	61023: The server has no GSS-TSIG key.
61024	Multistatus	61024: Server management via SNMP can only be in read-only. We recommend to update to SSL.
61025	Multistatus	61025: This server type is no longer supported. To manage a Microsoft server, create a "Microsoft DNS" server.
61026	Multistatus	61026: The zone has records configured with geolocation routing policy. You cannot edit or delete it from our GUI.
61027	Multistatus	61027: The zone has records configured with routing policy or health check option. You cannot edit or delete it from our GUI.

Appendix C. Default Gadgets

SOLIDserver offers gadgets by default in the *Gadgets Library*. Among them, 12 are displayed by default on the *Main Dashboard*, *NetChange dashboard* and *Device Manager dashboard*.

Table C.1. The default gadgets in the Gadgets Library

Gadget	Description	
Alerts	Displayed	No.
	Type	Top list.
	Description	The 10 latest raised alerts. Click on the alert name to go to the page <i>Alerts</i> . For more details, refer to the chapter Managing Alerts .
All networks	Displayed	No.
	Type	Quick Search.
	Description	A search engine gadget to find networks in the IPAM. The data entered in the fields automatically create filters on the page <i>All networks</i> through the columns <i>Name</i> and <i>Address</i> .
DHCP ranges	Displayed	No.
	Type	Top list.
	Description	A list of the 10 first DHCP ranges sorted by address with the server they belong to, their occupation rate and status.
DHCP Servers	Displayed	No.
	Type	Top list.
	Description	A list of the 10 first DHCP servers sorted by name with their type, protocol, synchronization status and status.
DNS RR type	Displayed	No.
	Type	Chart.
	Description	A pie chart representing the managed records repartition per type.
DNS Servers	Displayed	No.
	Type	Top list.
	Description	The 10 first DNS servers sorted by name with their IP address, synchronization status and status.
Alerts on ports/interfaces reconciliation drift	Displayed	On <i>Device Manager dashboard</i> .
	Type	Top list.
	Description	The 5 first port and/or interfaces in Device Manager that have a <i>Drift</i> in the Reconciliation column of the page <i>All ports & interfaces</i> . For more details, refer to the section Tracking Changes on the Page All ports & interfaces .
General information	Displayed	On the <i>Main Dashboard</i> .
	Type	Descriptive.
	Description	A gadget gathering the appliance configuration information: running services, hostname, IP address(es), default gateway(s), HA role and status. For more details regarding high availability, refer to the chapter Centralized Management .
Number of NetChange ports per device	Displayed	On <i>NetChange dashboard</i> .
	Type	Chart.
	Description	A pie chart representing NetChange ports repartition per network device.

Default Gadgets

Gadget	Description	
NetChange network devices vendor	Displayed	On <i>NetChange dashboard</i> .
	Type	Chart.
	Description	A pie chart representing NetChange network devices repartition per vendor.
NetChange active ports speed (bps)	Displayed	On <i>NetChange dashboard</i> .
	Type	Chart.
	Description	A pie chart representing NetChange <i>Active</i> ports repartition per speed. The <i>Inactive</i> and <i>Disabled</i> ports are not represented.
NetChange port status	Displayed	On <i>NetChange dashboard</i> .
	Type	Chart.
	Description	A pie chart representing NetChange ports repartition per status.
Number of interfaces used per device	Displayed	On <i>Device Manager dashboard</i> .
	Type	Chart.
	Description	A pie chart representing Device Manager used ports repartition per device.
Number of ports used per device	Displayed	On <i>Device Manager dashboard</i> .
	Type	Chart.
	Description	A pie chart representing Device Manager used interfaces repartition per device.
Number of RIPE assigned networks allocated per year	Displayed	No.
	Type	Chart.
	Description	A pie chart representing the RIPE assigned networks repartition per year.
Sum of allocated RIPE assigned networks size per year	Displayed	No.
	Type	Chart.
	Description	A pie chart representing the RIPE allocated networks repartition per size over the years.
Shortcuts	Displayed	On the <i>Main Dashboard</i> .
	Type	Shortcut.
	Description	A gadget containing shortcuts toward the IPAM, DNS and DHCP most commonly used pages.
Networks	Displayed	No.
	Type	Top list.
	Description	The 10 first networks sorted by start IP address with their start address and prefix, name, proportion of used IP addresses and status.
System Information	Displayed	On the <i>Main Dashboard</i> .
	Type	Descriptive.
	Description	A gadget gathering the appliance basic information: user connected, version, time and date, license type and expiration, manufacturer and product.
My account preferences & configuration	Displayed	On the <i>Main Dashboard</i> .
	Type	Shortcut.
	Description	A gadget gathering user dedicated shortcuts, toward the <i>Gadgets Library</i> and the wizard <i>Change Language</i> .
SOLIDserver Configuration Checklist	Displayed	On the <i>Main Dashboard</i> .
	Type	Configuration.
	Description	A gadget gathering appliance configuration shortcuts.

Appendix D. DHCP Options

Table of Contents

Basic Options	1245
Server Parameters	1246
Lease Information Options	1247
WINS/NetBIOS Options	1247
Host IP Options	1247
Interface Options	1248
Servers Options	1249
BOOTP Compatibility Options	1250
DHCP Packet Fields Options	1251
Microsoft DHCP Client Options	1252
NetWare Client Options	1252
NIS/NISplus Options	1253
Miscellaneous	1254
Vendor MSFT Options	1254
Vendor Nwip Options	1254

This appendix describes all the DHCP options that you can configure from the wizard **Configure DHCP options** at server, group, scope range and static level.

From either level, the wizard is accessible from the resource properties page, via the button **EDIT** of the panel *DHCP options*.

The options are distributed following the *Option category* of the wizard.

Basic Options

Table D.1. Basic DHCP options

Name	Code	Value type	Description
broadcast address	28	IP address	Specifies the broadcast address for the interface subnet.
domain name	15	text (name)	Domain name the client uses when resolving name via DNS.
domain-name-servers	6	list of IP addresses	List of Domain Name Servers (DNS) available for this client. These servers are listed by order of preference.
host name	12	text (name)	Client host name.
routers	3	list of IP addresses	List of routers for client subnet. These servers are listed by order of preference.
Authoritative	N/A	boolean	Allocation and checking of IP addresses according to the network segment where the DHCP client is connected.
Default lease time	N/A	duration (in seconds)	Default lease duration.
Max lease time	N/A	duration (in seconds)	Maximum lease duration (unavailable for BOOTP lease).
Min lease time	N/A	duration (in seconds)	Minimum lease duration.

Name	Code	Value type	Description
Ping check	N/A	boolean	Permit to check by an ICMP request if the target address is not used.
Ping time out	N/A	duration (in seconds)	Maximum timeout answer for a ping from the DHCP server.
Vendor option space	N/A	text	Define specific option space used for encapsulated options.
Site option space	N/A	text (name)	Permit to specify from which option space the <i>site-local</i> options are taken. The <i>site-local</i> options are the ones with an option code between 224 and 254.
Subnet mask	1	IP address	The subnet mask of the connected interface.

Server Parameters

These options concern the technical parameters on the server side.

Table D.2. Available server parameters

Name	Code	Value type	Description
Authoritative	N/A	boolean	Allocation and checking of IP addresses according to the network segment where the DHCP client is connected.
leasequery	N/A	boolean	Permit to specify if the server should respond to DHCPLEASEQUERY packets sent by CMTSS, i.e. send back lease information (creation/expiration date...).
Ping check	N/A	boolean	Permit to check by an ICMP request if the target address is not used.
Ping timeout	N/A	duration (in seconds)	Maximum timeout answer for a ping from the DHCP server.
Storm detection check request	N/A	number	Specifies the number of requests that have to be received in order to trigger the MAC address black listing. Only MAC addresses associated with an IP address are taken in account in the black list. It means that the clients have to make a DHCP request first with an IP address. This has to be used in conjunction with 'storm detection check sec' as well as 'Storm detection ignore sec' parameters. The value has to be between 1 and 65535.
Storm detection check sec	N/A	duration (in seconds)	Specifies the period during which the system allows requests. It then checks if it has more than X requests in this time lap, then if it is over, it blacklists the MAC for X seconds. It has to be used in conjunction with 'storm detection check request' as well as 'Storm detection ignore sec' parameters. The value has to be between 1 and 65535.
Storm detection check ignore sec	N/A	duration (in seconds)	Number of seconds during which any DHCP request from the blacklisted device should be ignored. Has to be used in conjunction with 'storm detection check request' as well as 'Storm detection check sec' parameters. The value has to be between 1 and 65535.

Lease Information Options

These options concern the technical mechanisms on the client side of SOLIDserver DHCP protocol.

Table D.3. Lease information options

Name	Code	Value type	Description
dhcp-renewal-time	58	duration (in seconds)	Time interval from address assignment until the client transitions to the RENEWING state.
dhcp-rebinding-time	59	duration (in seconds)	Time interval from address assignment until the client transitions to the REBINDING state.

WINS/NetBIOS Options

Table D.4. WINS/NetBIOS options

Name	Code	Value type	Description
netbios-name-servers	44	list of IP addresses	List of WINS servers or of Net-BIOS name servers (NBMS). These servers are sorted by order of preference. For more details, refer to <i>RFC1001</i> available on IETF website at http://tools.ietf.org/html/rfc1001 and to <i>RFC1002</i> at http://tools.ietf.org/html/rfc1002 .
netbios-dd-server	45	list of IP addresses	List of NetBIOS datagram distribution servers (NBDD), defined by <i>RFC1001</i> and <i>RFC1002</i> . These servers are sorted by order of preference.
netbios-node-type	46	number	Type of NetBIOS knot described in <i>RFC1001</i> and <i>RFC1002</i> . The value is represented by a numerical code: 1 for B-node, 2 for P-node, 4 for M-node, 8 for H-node.
netbios-scope	47	text (name)	Netbios-scope name value of NetBIOS scope specified in <i>RFC1001</i> and <i>RFC1002</i> .

Host IP Options

Table D.5. Host IP options

Name	Code	Value type	Description
Default-ip-ttl	23	duration (in seconds)	Default lifetime that the client must use to send a datagram on the network. Valid values between 1 and 255.
ip-forwarding	19	boolean	This option specifies whether the client should configure its IP layer for packets forwarding. For more details, refer to <i>RFC1533</i> available on IETF website at http://tools.ietf.org/html/rfc1533 .
Max-dgram-reassembly	22	number	Maximum size of datagram which the client must prepare to assemble.
non-local-source-routing	20	boolean	Allow the source-routing forwarding if the next-hop is on a different physical interface from that crossed by the datagram. For more details, refer to <i>RFC1122</i> available on IETF website at http://tools.ietf.org/html/rfc1122 .
path-mtu-aging-timeout	24	second	Aging time for the Path MTU Discovery defined for the client.

Name	Code	Value type	Description
			For more details, refer to <i>RFC1191</i> available on IETF website at http://tools.ietf.org/html/rfc1191 .
path-mtu-plateau-table	25	list of numbers	List of MTU sizes for the PMTU <i>RFC1191</i> . MTU sizes are prioritized by the order and do not have to be lower than 68.
policy-filter	21	2 IP addresses	Specifies the filtering policy for the non-local-source-routing. These filters are defined by a list of destination and netmask IP address couplets which specify the destination of entering routes. Any "routedsource" datagram not figuring in the list of filters is destroyed.
Subnet selection	118	IP address	The DHCP server determines the subnet from which the request originated. For more details, refer to <i>RFC3011</i> available on IETF website at http://tools.ietf.org/html/rfc3011 .

Interface Options

Table D.6. Interface options

Name	Code	Value type	Description
All-subnets-local	27	boolean	Specifies if the IP interface must demand that all subnets with which it communicates use the same MTU as that used by the physical interface.
Arp-cache-timeout	35	duration (in seconds)	This option specifies the timeout in seconds for ARP cache entries.
Auto configure	116	boolean	This option code is used to ask whether, and be notified if, auto-configuration should be disabled on the local subnet.
Broadcast-address	28	IP address	Specifies the broadcast address for the interface's subnet.
Classless static route	121	list of IP addresses	This option allows to use the routers used by the IP protocol to set up a packet transmission path between two IP hosts (one source and one destination host) through the router IP address, listed in the routing table. This option obsoletes the Static Route option (option 33). For more details, refer to the <i>RFC3442</i> available on IETF website at http://tools.ietf.org/html/rfc3442 .
Default-tcp-ttl	37	duration (in seconds)	This option specifies the default TTL that the client should use when sending TCP segments.
ieee802-3-encapsulation	36	boolean	Specifies if the client must use Ethernet Version 2 encapsulation or IEEE 802.3 on its interface if it is ethernet.
Interface-mtu	26	number	Size of MTU to use for this interface, it should be at least 68 bytes.
Mask-supplier	30	boolean	Specifies if the interface must declare its netmask during an ICMP echo.
Perform-mask-discovery	29	boolean	Specifies if, for this interface, the client should attempt an ICMP discovery to find its net-mask. Note that using this parameter is not recommended, as the first response received is taken into account and may be incorrect.

Name	Code	Value type	Description
Router-discovery	31	boolean	Specifies if, for this interface, the client should solicit routers by the "Router Discovery" mechanism. For more details, refer to the <i>RFC1256</i> available on IETF website at http://tools.ietf.org/html/rfc1256 .
Router-solicitation-address	32	IP address	Specifies the address by which, for this interface, the client must emit its solicitation requests to the routers.
Static-routes	33	2 IP addresses	In the route interface's cache, the first entry in the list is the destination address and the second is the router's address. The default route (0.0.0.0) is not tolerated here. This option was introduced in <i>RFC2132</i> but was obsoleted by the Classless Static Route Option (option 121).
Subnet-mask	1	IP address	The subnet mask for the network segment to which the client is connected.
Tcp-keepalive garbage	39	boolean	Specifies if the client must send a garbage byte with a <i>keepalive</i> message.
Tcp-keepalive-interval	38	duration (in seconds)	The time to wait before sending a keep alive message on a TCP connection.
Trailer-encapsulation	34	boolean	Specifies if the client must negotiate the use of trailers with ARP. For more details, refer to the <i>RFC893</i> available on IETF website at http://tools.ietf.org/html/rfc893 .

Servers Options

Table D.7. Server options

Name	Code	Value type	Description
Cookie-servers	8	list of IP addresses	Lists the cookie servers available for this client. These servers are listed by order of preference. For more details, refer to the <i>RFC865</i> available on IETF website at http://tools.ietf.org/html/rfc865 .
Finger-servers	73	list of IP addresses	List of Finger servers. These servers are sorted by order of preference.
Font-servers	48	list of IP addresses	Lists the system-X Windows font servers available for this client. These servers are sorted by order of preference.
len116-name-servers	5	list of IP addresses	IEN 116 name servers list for this client. These servers must be sorted by preference order.
Impress-server	10	list of IP addresses	Lists the Imagen Impress servers available for this client. These servers are listed by order of preference.
Irc-servers	74	list of IP addresses	List of Internet Relay Chat server.
Log-servers	7	list of IP addresses	Lists the UDP log servers (MIT-LCS syslog), available for this client. These servers are listed by order of preference.
Lpr-servers	9	list of IP addresses	Lists the printer servers available for this client. These servers are listed by order of preference. For more details, refer to the <i>RFC1179</i> available on IETF website at http://tools.ietf.org/html/rfc1179 .

DHCP Options

Name	Code	Value type	Description
Mobile-ip-home-agent	68	list of IP addresses	List the mobile IP home agent.
Nis-servers	41	list of IP addresses	Lists the IP of NIS servers available for the client. The servers can be sorted by order of preference.
Nis-plus-servers	65	list of IP addresses	Lists the IP addresses of NIS+ servers available for the client. The servers can be sorted by order of preference.
Ntp-servers	42	list of IP addresses	Lists the NTP news servers. These servers are sorted by order of preference.
Nntp-servers	71	list of IP addresses	Lists the NNTP news servers. These servers are sorted by order of preference.
Pop3-servers	70	list of IP addresses	Lists the POP3 message servers. These servers are sorted by order of preference.
Resource-location-servers	11	list of IP addresses	Lists the resource servers available for this client. These servers are listed by order of preference. For more details, refer to the <i>RFC887</i> available on IETF website at http://tools.ietf.org/html/rfc887 .
Smtpt-servers	69	list of IP addresses	Lists the SMTP message servers. These servers are sorted by order of preference.
Streetwork directory assistance server	76	list of IP addresses	Lists the IP addresses in order of preference for STDA servers available to the client.
Street-talk-servers	75	list of IP addresses	Lists the StreetTalk servers. These servers are sorted by order of preference.
Tftp-server-name	66	list of IP addresses	Name of the TFTP server to use when the Sname field is used to carry Options.
www-servers	72	list of IP addresses	Lists the WEB servers.
X-display-manager	49	list of IP addresses	Lists the X Window XDM system servers. These servers are sorted by order of preference.

BOOTP Compatibility Options

Table D.8. BOOTP compatibility options

Name	Code	Value type	Description
Boot-size	13	number	Length in block of 512 bytes of the boot image file for this client.
Boot-filename	67	number	Name of the boot file to use when the File field is used to carry options.
Cookie-servers	8	list of IP addresses	List the Cookie servers available. These servers are sorted by order of preference. For more details, refer to the <i>RFC865</i> available on IETF website at http://tools.ietf.org/html/rfc865 .
Domain-name-servers	6	list of IP addresses	Lists the domain name servers (DNS), available for this client. These servers are listed by order of preference.
Extensions-path	18	path	Name of the file containing additional options to be interpreted. The format is described in <i>RFC2132</i> . For more details, refer to the <i>RFC2132</i> available on IETF website at http://tools.ietf.org/html/rfc2132 .
Impress-server	10	list of IP addresses	Lists the Imagen Impress servers available for this client. These servers are listed by order of preference.

DHCP Options

Name	Code	Value type	Description
Merit-dump	14	path	Path of file in which the client must copy the memory image in the event of a crash. This path is constituted by a set of NVT ASCII characters.
Resource-location-servers	11	list of IP addresses	Lists the resource servers available for this client. These servers are listed by order of preference. For more details, refer to the <i>RFC887</i> available on IETF website at http://tools.ietf.org/html/rfc887 .
Root-path	17	path	Path of the disk route for this client. This path is constituted by a set of NVT ASCII characters.
Filename	N/A	file name	Name of the boot file to use when the field is used to carry options.
Next-server	N/A	IP address	This option allows to specify the IP address of the server from which the initial boot file (specified in the filename statement) has to be loaded. Server-name should be a numeric IP address. If no next-server parameter applies to a given client, the DHCP server's IP address is used. Some clients prefer to receive the server name in the server-name option.
Server-name	N/A	text (name)	this statement can be used to inform the client of the name of the server from which it is booting. This name should be the same as the one provided to the client.
Swap-server	16	IP address	Swap server.
Time-offset	2	duration (in seconds)	Time offset from UTC (Coordinated Universal Time).
Time-servers	4	IP address	Time server available for this DHCP client.

DHCP Packet Fields Options

Table D.9. Packet fields options

Name	Code	Value type	Description
DHCP client identifier	61	text	For client that want uses the different identifier, DHCP defines the <i>client identifier</i> option. This option tells the server to use the value in the option to identify the client, rather than using the client MAC address.
DHCP parameter request list	55	list of numbers	Used by a DHCP client to request specific option type values from the DHCP server. Each option type is requested and listed by a number value containing a valid or recognized DHCP option code for the server.
Dhcp-rebinding-time	59	duration (in seconds)	Specifies the time interval from address assignment until the client transitions to the REBINDING state.
Dhcp-renewal-time	58	duration (in seconds)	Specifies the time interval from address assignment until the client transitions to the RENEWING state.
Dhcp-server-identifier	54	IP address	The identifier is the IP address of the selected server.
User-class	77	text	Information on the client class.
Vendor-class-identifier	60	text	This option is used by DHCP clients to optionally identify the vendor type and configuration of a DHCP client.
Vendor-encapsulated-options	63	<i>provided by the vendor</i>	This option allows to use encapsulated options provided by your vendor. The option can contain a single vendor-specific value or one (or more) vendor-specific sub-options. The configuration set for each vendor is saved in a class. All the vendor classes are

Name	Code	Value type	Description
			kept separately in the DHCP server configuration and trigger a unique response for each vendor.

Microsoft DHCP Client Options

Table D.10. Microsoft DHCP client options

Name	Code	Value type	Description
dhcp-lease-time	51	duration (in seconds)	This option is used in a client request (<i>DHCPDISCOVER</i> or <i>DHCPREQUEST</i>) to allow the client to request a lease time for the IP address.
dhcp-rebinding-time	59	duration (in seconds)	Specifies the time interval from address assignment until the client transitions to the REBINDING state.
dhcp-renewal-time	58	duration (in seconds)	Specifies the time interval from address assignment until the client transitions to the RENEWING state.
dhcp-server-identifier	54	address	The identifier is the IP address of the selected server.
domain name	15	name	Specifies the domain name that client should use when resolving hostnames via the Domain Name System.
domain-name-servers	6	list of IP addresses	Specifies a list of Domain Name System name servers available to the client. Servers should be listed in order of preference.
Domain search list	135	list of domains	In some circumstances, it is useful for the DHCP client to be configured with the domain search list. Note that Microsoft Windows 200x, XP do not support a list of domain search.
netbios-name-servers	44	list of IP addresses	List of Net-BIOS name servers (NBMS). These servers are sorted by order of preference. For more details, refer to <i>RFC1001</i> available on IETF website at http://tools.ietf.org/html/rfc1001 and to <i>RFC1002</i> at http://tools.ietf.org/html/rfc1002 .
netbios-node-type	46	hexadecimal	The NetBIOS node type option allows NetBIOS over TCP/IP clients which are configurable to be configured as described in <i>RFC1001</i> and <i>RFC1002</i> . Available values are: 0x1 = B-node; 0x2 = P-node; 0x4 = M-node; 0x8 = H-node
netbios-scope	47	name	Specifies the NetBIOS over TCP/IP scope parameter for the client as specified in <i>RFC1001</i> and <i>RFC1002</i> .
Routers	3	list of IP addresses	Specifies a list of IP addresses for routers on the client's subnet. Routers should be listed in order of preference.
WWW proxy server	252	URL	This option is used to automatically configure proxy settings for the client's browser. Type in the URL of the server that stores the information.

NetWare Client Options

Table D.11. NetWare client options

Name	Code	Value type	Description
Nds-context	87	text	Specifies the initial NDS context the client should use.

Name	Code	Value type	Description
Nds-servers	85	IP address	Specifies one or more NDS servers for the client to contact for access to the NDS database. Servers should be listed in order of preference.
Nds-tree-name	86	name	Specifies the initial NDS context the client should use.
Nwip-domain	62	name	This option code is used to convey the NetWare/IP domain name used by the NetWare/IP product.
Slp-directory-agent	78	address IP	Specifies the location of one or more SLP Directory Agents.
Slp-service-scope	79	scope	Indicates the scopes that a SLP Agent is configured to use.

NIS/NISplus Options

Table D.12. NIS/NISplus options

Name	Code	Value type	Description
Nis-domain	40	name	Specifies the name of the client's NIS domain. The domain is formatted as a character string consisting of characters from the NVT ASCII character set.
Nis-servers	41	list of IP addresses	Lists the IP of NIS servers available for the client. The servers can be sorted by order of preference.
Nis-plus-domain	64	name	Specifies the name of the client's NIS+ domain. The domain is formatted as a character string consisting of characters from the NVT ASCII character set.
Nis-plus-servers	65	list of IP addresses	Specifies a list of IP addresses indicating NIS+ servers available to the client. Servers should be listed in order of preference.
Autoretires	8	<i>provided by the vendor</i>	This option specifies a list of Quote of the Day servers available to the client. The servers SHOULD be listed in order of preference.
Autoretry secs	9	<i>provided by the vendor</i>	This option specifies a list of LPR servers available to the client. The servers SHOULD be listed in order of preference.
Nearest nwip server	7	<i>provided by the vendor</i>	This option specifies a list of MIT-LCS UDP servers available to the client. The servers SHOULD be listed in order of preference.
Nsq broadcast	5	<i>provided by the vendor</i>	This option specifies a list of Name servers available to the client. The servers SHOULD be listed in order of preference.
Nwip 1 1	10	<i>provided by the vendor</i>	This option specifies a list of Imagen Impress servers available to the client. The servers SHOULD be listed in order of preference.
Preferred dss	6	<i>provided by the vendor</i>	This option specifies a list of DNS servers available to the client. The servers SHOULD be listed in order of preference.
Primary dss	11	<i>provided by the vendor</i>	This option specifies a list of RLP servers available to the client. The servers SHOULD be listed in order of preference.

Miscellaneous

Table D.13. Other DHCP options

Name	Code	Value type	Description
Domain search	119	list of domains	DNS domain search list. For more details, refer to the <i>RFC3397</i> available on IETF website at https://tools.ietf.org/html/rfc3397 .
Name service search	117		Name Service Search.

Vendor MSFT Options

Table D.14. Vendor MSFT options

Name	Code	Value type	Description
Default routers TTL	3	list of IP addresses	This option specifies a list of 32 bit IP addresses for routers on the client's subnet. The routers SHOULD be listed in order of preference.
Disable netbios	1	<i>provided by the vendor</i>	The subnet mask for the network segment to which the client is connected.
Release on shutdown	2	<i>provided by the vendor</i>	This option specifies the offset of the client's subnet in seconds from Coordinated Universal Time (UTC).

Vendor Nwip Options

Table D.15. Vendor Nwip options

Name	Code	Value type	Description
Autoretries	8	<i>provided by the vendor</i>	This option specifies a list of Quote of the Day servers available to the client. The servers SHOULD be listed in order of preference.
Autoretry secs	9	<i>provided by the vendor</i>	This option specifies a list of LPR servers available to the client. The servers SHOULD be listed in order of preference.
Nearest nwip server	7	<i>provided by the vendor</i>	This option specifies a list of MIT-LCS UDP servers available to the client. The servers SHOULD be listed in order of preference.
Nsq broadcast	5	<i>provided by the vendor</i>	This option specifies a list of Name servers available to the client. The servers SHOULD be listed in order of preference.
Nwip 1 1	10	<i>provided by the vendor</i>	This option specifies a list of Imagen Impress servers available to the client. The servers SHOULD be listed in order of preference.
Preferred dss	6	<i>provided by the vendor</i>	This option specifies a list of DNS servers available to the client. The servers SHOULD be listed in order of preference.
Primary dss	11	<i>provided by the vendor</i>	This option specifies a list of RLP servers available to the client. The servers SHOULD be listed in order of preference.

Appendix E. MAC Address Types References

This appendix lists all the MAC address types used in SOLIDserver that you can display on the page *DHCP All statics* both in IPv4 and IPv6. There is a set of 31 different types of MAC addresses that you can specify when adding or editing DHCP statics. Each type corresponds to a protocol that has been assigned a reference number defined in the IANA Address Resolution Protocol (ARP). In the GUI, this reference adds an extra byte at the beginning of the MAC addresses listed in the column **MAC address** on the page *All statics*. Typically, the MAC addresses listed in this column look as follows: <1_byte_MAC_type_reference>:<6_bytes_MAC_address>.

The different types of MAC addresses can be listed separately from the MAC address itself using the *DHCP static MAC type* column. This column displays two columns: the column **MAC type** that displays the MAC type code (except for Ethernet that is listed in full letters) and the column **MAC address** that displays the MAC address in its traditional format.

Note that every **reference is listed in hexadecimal form** in the wizard. Therefore, the ARP parameter *10* (for Autonet) is listed as *0a* and so forth.

Table E.1. Supported MAC address types references

MAC type	Reference
Unknown	You can use any hexadecimal reference number, as long as it is not already listed below.
Ethernet	01
Experimental ethernet	02
Amateur radio AX25	03
Proteon ProNET Token Ring	04
Chaos	05
Token Ring	06
ARCNET	07
FDDI	08
Lanstar	09
Autonet	0a
LocalTalk	0b
LocalNet	0c
Ultralink	0d
SMDS	0e
Frame Relay	0f
ATM	15
HDLC	11
Fibre Channel	12
Serial Line	14
MIL-STD-188-220	16
Metricom	17
IEEE 1394.1995	18
MAPOS	19

MAC Address Types References

MAC type	Reference
Twin Axial	1a
EUI-64	1b
HIPARP	1c
IP/ARP over ISO 7816-3	1d
ARPSec	1e
IPSec tunnel	1f
InfiniBand	20

Appendix F. DNS Resource Records Related Fields

This appendix provides a record per record description of the different fields to configure when adding a resource record to a master zone. For more details regarding each record, refer to the section [Adding Resource Records](#).

To add a resource record to a zone

1. In the sidebar, go to **DNS > RRs**. The page **All RRs** opens.
2. In the column **Zone**, click on the name of the zone of your choice to display its records.
3. In the menu, click on **+ Add**. The wizard **Add a DNS RR** opens.
4. In the drop-down list **RR type**, select the record of your choice.
5. Set the **RR Name** and **TTL** according to your needs.
6. Configure the record fields following the table [DNS resources records configuration fields](#).
7. Click on **OK** to complete the operation. The report opens and closes. The record is now listed.

Note that all the records supported are not supported by all DNS servers. For more details, refer to the section [Adding Resource Records](#).

Table F.1. DNS resources records configuration fields

RR type	Related field(s)	Description
A	IP address	The IPv4 Address of the host.
AAAA	IPv6 address	The IPv6 Address of the host.
AFSDB	Preference	The version of AFS service used: 1 (AFS version 3.0) or 2 (OSF DCE/NCA version).
	AFS server	The AFS hostname.
CAA	Flags	An integer, between 0 and 255, that influences the interpretation of the record. 0 means that the record is not critical.
	Property identifier	The CAA record tag, including but not limited to: <i>issue</i> authorizes a single certificate authority to issue a certificate of any type for the hostname.
		<i>issuewild</i> authorizes a single certificate authority to issue a wildcard certificate for the hostname.
		<i>iodef</i> specifies the URL to which a certificate authority may report policy violations.
Value	The domain of the CA associated with the tag or a URL, depending on the tag specified in the <i>Property identifier</i> .	
CERT	Type	An integer 0 and 65535 that specifies the type of certificate where 1 is PKIX, 2 is SPKI, 3 is PGP, 4 is IPKIX, 5 is ISPKI, 6 is IPGP, 7 is ACPKIX, 8 is IACPKIX, 253 is URI and 254 is OID.
	Key tag	The certificate's key tag, a 16-bit value computed for the key embedded in the certificate.
	Algorithm	The public key's cryptographic algorithm.
	Certificate or CRL	The certificate, encoded in base-64 format.
	<i>The record CERT is described in RFC 4398 available at https://tools.ietf.org/html/rfc4398.</i>	

DNS Resource Records Related Fields

RR type	Related field(s)	Description
CNAME	Hostname	The hostname.
DHCID	Key	The encoded DHCP client identifier (DUID) in a binary format, encoded in base-64.
	<i>For more details, refer to the section 3 of RFC 4701, available at https://tools.ietf.org/html/rfc4701#section-3.</i>	
DNAME	Domain	The name of a subdomain of the zone.
HINFO	CPU	Select the CPU description in the drop-down list. If yours is not listed, type it in the field and let the default value in the list (<i>Other</i>).
	OS	Select the OS in the drop-down list. If yours is not listed, type it in the field and let the default value in the list (<i>Other</i>).
MINFO	Responsible email	The email address of the administrator of the mail list.
	Error email	The email address that should receive the error messages regarding the mail list.
MX	Preference	An integer, between 0 and 65535, that defines which server has priority if there are several records in the zone. The lowest the value has the priority over the other server(s).
	Mail server	The mail server hostname.
	<i>The accepted configuration of MX and NS records is detailed in the section 10.3 of RFC 2181, available at http://tools.ietf.org/html/rfc2181.</i>	
NAPTR	Order	An integer, between 0 and 65535, that defines which RR has priority if there are several NAPTR records in the zone. The lowest the value has the priority over the other record(s).
	Preference	An integer, between 0 and 65535, that defines which RR has priority if there are several NAPTR records have the same order in the zone. The lowest the value has the priority over the other record(s).
	Flags	The string that corresponds to the action you want your client application to perform. The flag specified impacts the data expected in the field <i>Services</i> , <i>Regex</i> and/or <i>Replace</i> .
	Services	The services parameters to which applies the action specified in the field <i>Flags</i> . You must respect your client application syntax.
	Regex	The string that contains a substitution expression matching the format <code><delimit ereg delimit substitution delimit flag></code> to which applies the action specified the field <i>Flags</i> .
	Replace	An FQDN domain name to which applies the action specified the field <i>Flags</i> . You can specify no domain name if you type in . (dot) in the field.
	<i>The record NAPTR is described in RFC 3403, available at http://tools.ietf.org/html/rfc3403.</i>	
NS	DNS server	The DNS server hostname.
	<i>The accepted configuration of MX and NS records is detailed in the section 10.3 of RFC 2181, available at http://tools.ietf.org/html/rfc2181.</i>	
NSAP	Name	The NSAP address of the end system. It should start with 0x and not exceed 255 hexadecimal characters separated by dots.
	<i>The NSAP RR is described in the RFC 1706, available at http://tools.ietf.org/html/rfc1706.</i>	
OPENPGPKEY	Key	The OpenPGP public key, without the last line. The last line is preceded by the character =.
	<i>The record OPENPGPKEY is described in RFC 7629, available at https://tools.ietf.org/html/rfc7629.</i>	
PTR	Localization	The hostname that should be returned when the address is queried.
SSHFP	Algorithm	An integer that identifies the algorithm used by the Public Key: 1 (RSA), 2 (DSA), 3 (ECDSA) or 4 (Ed25519).

DNS Resource Records Related
Fields

RR type	Related field(s)	Description
	Type	An integer that identifies the type of fingerprint used: 1 (SHA-1) or 2 (SHA-256).
	Fingerprint	An hexadecimal string representing the hash result.
SRV	Priority	An integer, between 0 and 65535, that defines which server has priority if there are several SRV records in the zone. The lowest the value has the priority over the other server(s).
	Weight	An integer, between 0 and 65535, that defines the server weight. If two SRV records have the same priority, the weight defines which one is more used. The field gives priority to the SRV RR with the greatest weight value: the greater the value is, the more the server is solicited. If you type in 0, there is no weighting.
	Ports	The port number that delivers the service to the target.
	Target	The hostname of the server delivering the service.
TXT	Text	The description of your choice (max. 255 characters including spaces).
TLSA	Certificate Usage	The association data provided to match the certificate presented in the TLS handshake. It is identified via an integer, between 0 and 255, among them:
		1 specifies an end entity certificate, or the public key the certificate. It must be matched with the end entity certificate given by the server in TLS.
		2 specifies a certificate, or the public key the certificate. It must be used as the Trust anchor when validating the end entity certificate given by the server in TLS.
		3 specifies a certificate, or the public key the certificate. It must match the end entity certificate given by the server in TLS.
	Selector	The part of the TLS certificate that the server presents to compare with the association data, where 0 is the Full certificate and 1 is the SubjectPublicKeyInfo.
	Matching Type	The way the association data is presented, where 0 is the Exact match on selected content, 1 is an SHA-256 hash of selected content and 2 is an SHA-512 hash of selected content.
	Certificate Association Date	The <i>certificate association data</i> to be matched. The expected value in the field is a string of hexadecimal characters, it can include spaces.
<i>The record TLSA is described in RFC 6698, available at https://tools.ietf.org/html/rfc6698.</i>		
URI	Priority	An integer, between 0 and 65535, that defines which target URI has priority if there are several URI records in the zone. The lowest value has the priority over the other URI.
	Weight	An integer, between 0 and 65535, that defines the target URI weight. If two URI records have the same priority, the weight defines which one is more used. The field gives priority to the URI record with the greatest weight value: the greater the value is, the more the URI is solicited. If you type in 0, there is no weighting.
	Target	The targeted URI, specified using the syntax described in the RFC3986.
	<i>The record URI is described in RFC 3986, available at http://tools.ietf.org/html/rfc3986.</i>	
WKS ^a	IP address	The IPv4 Address of the host that contains the services listed in the Services field.
	Protocol	The protocol that suites your needs.
	Services	The list of needed services.
DNSSEC records		

DNS Resource Records Related
Fields

RR type	Related field(s)	Description
CDNSKEY	Flags	The zone key flag.
	Protocol	The protocol value.
	Algorithm	The public key's cryptographic algorithm.
	Key	The public key material.
	<i>The records CDNSKEY is described in RFC 7344, available at http://tools.ietf.org/html/rfc7344.</i>	
CDS	Key Tag	The key tag of the DS record of the zone used in the delegation.
	Key Algorithm	The algorithm key of the DS record of the zone used in the delegation.
	Digest Type	The digest type of the DS record of the zone used in the delegation.
	Digest	The digest of the DS record of the zone used in the delegation.
	<i>The records CDS is described in RFC 7344, available at http://tools.ietf.org/html/rfc7344.</i>	
DNSKEY	Flags	The zone key flag.
	Protocol	The protocol value.
	Algorithm	The public key's cryptographic algorithm.
	Key	The public key material.
DS	Key Tag	The child zone DS key tag.
	Key Algorithm	The child zone DS algorithm key.
	Digest Type	The child zone DS digest type.
	Digest	The child zone DS digest.

^aThis record type is obsolete.

Appendix G. Advanced Properties

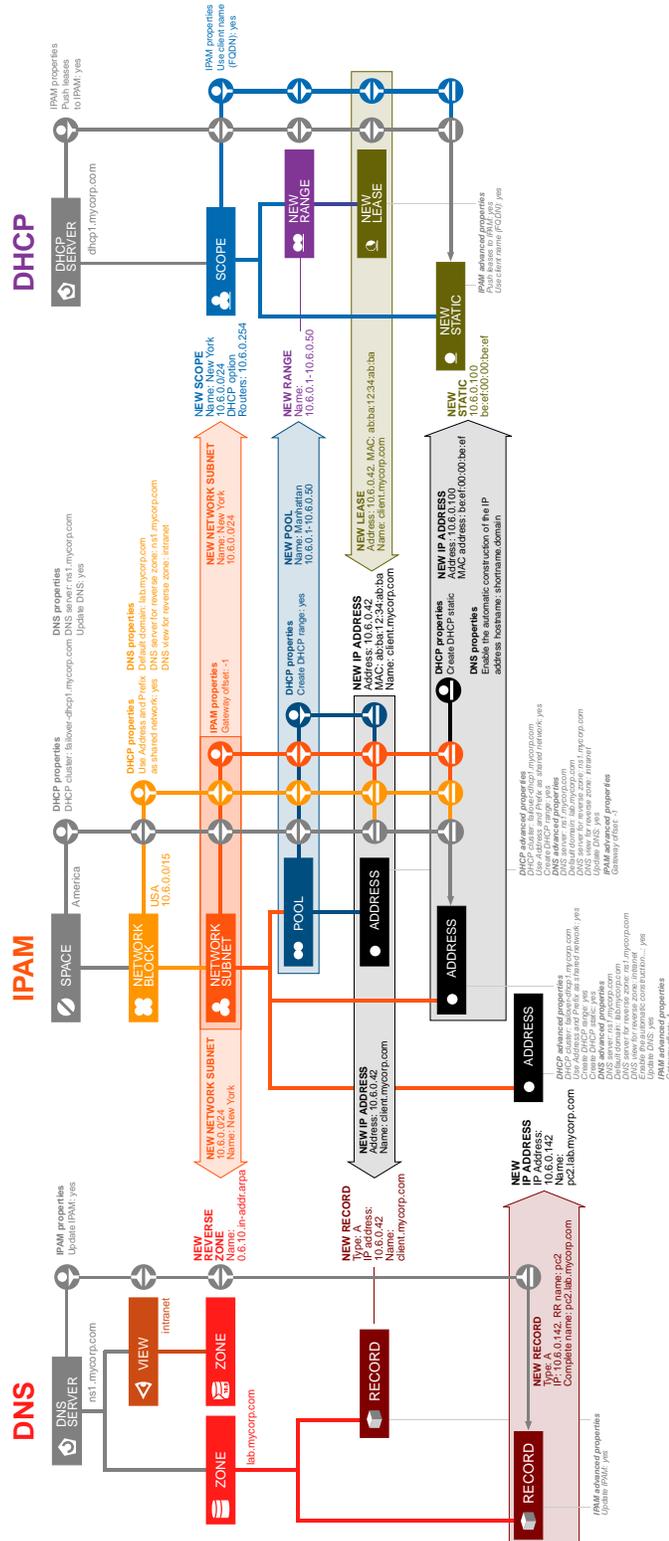


Figure G.1. All the advanced properties

Appendix H. User Tracking Services Filter

This appendix provides a list of the available filters of the drop-down list *Services* on the page User Tracking. For more details regarding this page, refer to the section [Tracking Users](#).

Table H.1. The available filters of the drop-down list Services

Service	Description
DHCP	All the DHCP services
DHCP server	All the DHCP server related operations.
Add: DHCP servers	All the DHCP server additions and editions.
Add: DHCPv6 servers	All the DHCPv6 server additions and editions.
Delete: DHCP servers	All the DHCP server deletions.
Delete: DHCPv6 servers	All the DHCPv6 server deletions.
DHCP scope	All the DHCP scope related operations.
Add: DHCP scopes	All the DHCP scope additions and editions.
Add: DHCPv6 scopes	All the DHCPv6 scope additions and editions.
Delete: DHCP scopes	All the DHCP scope deletions.
Delete: DHCPv6 scopes	All the DHCPv6 scope deletions.
DHCP range	All the DHCP range related operations.
Add: DHCP ranges	All the DHCP range additions and editions.
Add: DHCPv6 ranges	All the DHCPv6 range additions and editions.
Delete: DHCP ranges	All the DHCP range deletions.
Delete: DHCPv6 ranges	All the DHCPv6 range deletions.
DHCP static	All the DHCP static related operations.
Add: DHCP static	All the DHCP static additions and editions.
Add: DHCPv6 static	All the DHCPv6 static additions and editions.
Delete: DHCP statics	All the DHCP static deletions.
Delete: DHCPv6 statics	All the DHCPv6 static deletions.
DHCP option	All the DHCP option related operations.
Add: DHCP options	All the DHCP option additions and editions.
Add: DHCPv6 options	All the DHCPv6 option additions and editions.
Delete: DHCP options	All the DHCP option deletions.
Delete: DHCPv6 options	All the DHCPv6 option deletions.
IPAM	All the IPAM services
Space	All the IP space related operations.
Add: spaces	All the IP space additions and editions.
Delete: spaces	All the IP space deletions.
Network	All the network related operations.
Add: IPv4 networks	All the IPv4 network additions and editions.
Add: IPv6 networks	All the IPv6 network additions and editions.
Delete: IPv4 networks	All the IPv4 network deletions.
Delete: IPv6 networks	All the IPv6 network deletions.

User Tracking Services Filter

Service	Description
Pool	All the pool related operations.
Add: IPv4 pools	All the IPv4 pool additions and editions.
Add: IPv6 pools	All the IPv6 pool additions and editions.
Delete: IPv4 pools	All the IPv4 pool deletions.
Delete: IPv6 pools	All the IPv6 pool deletions.
Address	All the IP address related operations.
Add: IPv4 addresses	All the IPv4 address additions and editions.
Add: IPv6 addresses	All the IPv6 address additions and editions.
Delete: IPv4 addresses	All the IPv4 address deletions.
Delete: IPv6 addresses	All the IPv6 address deletions.
Alias	All the aliases related operations.
Add: aliases to IPv4 addresses	All the IPv4 alias additions and editions.
Add: aliases to IPv6 addresses	All the IPv6 alias additions and editions.
Delete: Pv4 addresses aliases	All the IPv4 alias deletions.
Delete: IPv6 addresses aliases	All the IPv6 alias deletions.
DNS	All the DNS services
DNS server	All the DNS server related operations.
Add: DNS servers	All the DNS server additions and editions.
Delete: DNS servers	All the DNS server deletions.
DNS zone	All the DNS zone related operations.
Add: DNS zones	All the DNS zone additions and editions.
Delete: DNS zones	All the DNS zone deletions.
DNS RR	All the DNS record related operations.
Add: DNS RRs	All the DNS record additions and editions.
Delete: DNS RRs	All the DNS record deletions.
Guardian	All the Guardian services
Policy	All the policy related operations.
Add: policies	All the policy additions and editions.
Delete: policies	All the policy deletions.
Trigger	All the trigger related operations.
Add: triggers	All the trigger additions and editions.
Delete: triggers	All the trigger deletions.
Guardian parameters	All the guardian parameter additions.
Application	All the Application related operations
Application	All the application related operations.
Add: applications	All the application additions and editions.
Delete: applications	All the application deletions.
Pool	All the pool related operations.
Add: pools	All the pool additions and editions.
Delete: pools	All the pool deletions.
Node	All the node related operations.
Add: nodes	All the node additions and editions.
Delete: nodes	All the node deletions.

User Tracking Services Filter

Service	Description
NetChange	All the NetChange related operations
Network device	All the network device related operations.
Add: network devices	All the network device additions.
Edit: network devices	All the network device editions.
Delete: network devices	All the network device deletions.
Edit: ports properties	All the port property editions.
VLAN	All the VLAN related operations.
Add: VLANs	All the VLAN additions.
Delete: VLANs	All the VLAN deletions.
Add: a VLAN to a port	All the additions of VLAN to a port.
Delete: a VLAN from a port	All the deletions of VLAN from a port.
Rule	All the rule related operations
Add: rules	All the rule additions and editions.
Delete: rules	All the rule deletions.
Group	All the group of users related operations
Add: groups	All the group of users additions and editions.
Delete: groups	All the group deletions.
Add: user as group resource	All the additions of users as resource of a group.
Remove: user from group resource	All the deletions of users from the resources of a group.
User	All the users related operations
Add: users	All the user additions and editions.
Delete: users	All the user deletions.
System	All the system related operations
Install: Packages	All the operations related to package installation.
Uninstall: Packages	All the operations related to package uninstallation.
Generate: SSL certificates	All the operations related to SSL certificates generation.
Delete: files uploaded to Local Files Listing	All the deletion operations of files uploaded to Local Files Listing.
Upload: files to Local Files Listing	All the upload operations of files to Local Files Listing.
Edit: Remote archive	All the remote archive editions.
License	All the license related operations
Add: license	All the license additions and editions.
Delete: license	All the license deletions.
Class	All the Class Studio related operations
Add: classes	All the class additions and editions.
Delete: classes	All the class deletions.
Gadgets	All the gadgets related operations
Add: gadgets to dashboards	All the additions of gadgets to dashboards.
Delete: gadgets from dashboards	All the deletions of gadgets from dashboards.
Bookmarks	All the bookmarks related operations
Add: bookmarks	All the bookmark additions and editions.
Delete: bookmarks	All the bookmark deletions.
Alert Definition	All the alert definition related operations
Add: alert definitions	All the alert definition additions and editions.

User Tracking Services Filter

Service	Description
Delete: alert definitions	All the alert definition deletions.
Workflow	All the workflow related operations
Add: requests	All the request additions and editions.
Delete: requests	All the request deletions.
Custom DB	All the custom DB related operations
Custom database	All the custom database related operations.
Add: custom databases	All the custom database additions and editions.
Delete: custom databases	All the custom database deletions.
Custom data	All the custom data related operations.
Add: data to a custom database	All the additions of data to a custom database.
Delete: data from a custom database	All the deletions of data from custom database.

Appendix I. SNMP Metrics

Table of Contents

Prerequisites	1266
Understanding the SNMP Metrics Presentation	1267
Retrieving SNMP Metrics via CLI	1267
Monitoring the Hardware	1268
Monitoring the System	1270
Monitoring the DHCP Service	1274
Monitoring the DNS Service	1275
Monitoring DNS Guardian	1278

This appendix provides a list of the most relevant indicators, the SNMP metrics, you can monitor from an external solution.

The SNMP metrics are detailed in Management Information Bases (MIB) that formally describe the network objects you can monitor. The MIB content is hierarchized thanks to a suite of numbers called Object IDentifiers (OID). All the OIDs of a MIB are organized like a tree of information with common trunks appended by unique ends that refer to a specific node, a unique set of information. Each OID node can therefore match a network object, such as the status of a power supply, or a specific property of the network object, like a variable name or values.

This appendix includes proprietary, IANA, IEEE and IETF managed MIBs:

- IDRAC-MIB
- UCD-SNMP-MIB
- HOST-RESOURCES-MIB
- IF-MIB
- EIP-STATS
- EIP-DNSGUARDIAN
- EIP-MON-MIB

Prerequisites

- Have an Internet connection and your credentials ready to download the files *.mib on our website, at <https://downloads.efficientip.com/support/downloads/MIBs/>. If you do not have credentials yet, request them at www.efficientip.com/support-access.

You can get a tree overview of each MIB at <http://www.oidview.com>.

- Make sure each monitored SOLIDserver is configured to allow the SNMP collector to retrieve the SNMP information:
 - From SOLIDserver, you must configure the SNMP agent with an access list using either community strings in SNMP v1/v2c, or authentication credentials in SNMP v3.

By default, a v1/v2c profile exists with the community string *public*. For more details, refer to the section [Managing the SNMP Service](#).

- From the SNMP collector's side, you must use SNMP v2c or v3 and configure the external monitoring solution to ensure it can access SOLIDserver SNMP agent and leverage the metrics that you need.

Note that you can also configure SNMP traps.

Once your system is properly configured, you can set various SNMP alerts on SOLIDserver objects to be notified of any unusual behavior. For more details, refer to the chapter [Managing Alerts](#).

Understanding the SNMP Metrics Presentation

Within this appendix, OIDs are detailed in tables as follows:

Table I.1. How OIDs are described

.X.X.X.X.X.XXX.XXXXX.X.X.XXX.XX.X.X The OID of the metric	
Variable name: <name>. The name of the OID, as declared in the MIB	Type: <type>
Description: <description>. Information on the OID.	The OID value format.
Expected values: <value>. The value(s) that the OID can return.	

In addition, the OIDs presented can be followed by an asterisk *.

The asterisk represents subsets, for which, the default value of the first object can be 0 or 1. In the case of the OID `1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.*`, if you have two power supply units:

- You can specify `1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.0` to retrieve the status of the first supply unit, or
- You can specify `1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1` to retrieve the status of the second one.

Retrieving SNMP Metrics via CLI

You can retrieve the value of an object manually via CLI with a command `snmpget` or `snmpwalk` on the monitored host IP address.

In the following example, we execute the command `snmpget` via SNMP v2c with the community string `public` to retrieve the OID `1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1` (the status of the first power supply) and the value 3 (the status ok). This OID is described the IDRAC-MIB:

```
snmpget -v2c -c public -On<host-IP-address>
.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1 = INTEGER: 3
```

You can also use the command `snmpwalk` to search and list information of all the objects in a particular subset. Here, requesting `1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1` returns the status of two power supplies:

```
snmpwalk -v2c -c public -On<host-IP-address>
.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.1 = INTEGER:
3.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.2 = INTEGER: 3
```

Note that for OIDs followed by an asterisk *:

- To execute an `snmpget` command on a subset, you must specify the number of the object you are requesting. The default value of the first object can be 0 or 1.
- To execute an `snmpwalk` command on a subset, specify the OID without the `.*` characters.

Monitoring the Hardware

Note that **the metrics in this section only apply to hardware appliances**. They provide information on physical hardware components to ensure they are perfectly running and maximize hosted service availability. To monitor virtual and software SOLIDserver appliances, refer to the other sections.

A Nagios compatible plug-in is available on Github: https://github.com/dangmocrang/check_idrac.

Power Unit

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.2. Power unit

.1.3.6.1.4.1.674.10892.5.4.600.10.1.8	
Variable name: powerUnitStatus	Type: integer
Description: The status of the power unit.	
Expected values: <i>ok</i> (i.e. 3)	

Power Redundancy

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.3. Power redundancy

.1.3.6.1.4.1.674.10892.5.4.600.10.1.5	
Variable name: powerUnitRedundancyStatus	Type: integer
Description: The redundancy status of the power unit.	
Expected values: <i>full</i> or <i>notRedundant</i> (i.e. 3 or 6)	

Power Supply

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.4. Power supply

.1.3.6.1.4.1.674.10892.5.4.600.12.1.5.1.*	
Variable name: powerSupplyStatus	Type: integer
Description: The status of the power supply.	
Expected value: <i>ok</i> (i.e. 3)	
.1.3.6.1.4.1.674.10892.5.4.600.12.1.9.1.*	
Variable name: PowerSupplyInputVoltage	Type: integer
Description: The current power supply input in Volts.	

CPU

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.5. CPU

.1.3.6.1.4.1.674.10892.5.4.600.50.1.5.1.*	
Variable name: ProcessorDeviceStatus	Type: integer
Description: The status of the CPU socket.	

Expected values: <i>ok</i> (i.e. 3)
--

Memory

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.6. Memory

<i>.1.3.6.1.4.1.674.10892.5.4.1100.50.1.5.1.*</i>	
Variable name: MemoryDeviceStatus	Type: integer
Description: The status of the memory slot.	
Expected values: <i>ok</i> (i.e. 3)	

Virtual Disk

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.7. Virtual disk

<i>.1.3.6.1.4.1.674.10892.5.5.1.20.140.1.1.4.*</i>	
Variable name: VirtualDiskState	Type: integer
Description: The status of the RAID Controller Status for the virtual disk.	
Expected values: <i>online</i> (i.e. 3)	

Physical Disk

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.8. Physical disk

<i>.1.3.6.1.4.1.674.10892.5.5.1.20.130.4.1.4.*</i>	
Variable name: PhysicalDiskState	Type: integer
Description: The status of the physical disk.	
Expected values: <i>online</i> (i.e. 3)	

Temperature

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.9. Temperature

<i>.1.3.6.1.4.1.674.10892.5.4.700.20.1.5.1.*</i>	
Variable name: TemperatureProbeStatus	Type: integer
Description: The status of the temperature probe.	
Expected values: <i>ok</i> (i.e. 3)	
<i>.1.3.6.1.4.1.674.10892.5.4.700.20.1.6.1.*</i>	
Variable name: TemperatureProbeReading	Type: integer
Description: The current temperature in tenths of degrees Centigrade.	

Fan

All the OIDs described in this section are part of the MIB extension **IDRAC-MIB**.

Table I.10. Fan

.1.3.6.1.4.1.674.10892.5.4.700.12.1.5.1.*	
Variable name: CoolingDeviceStatus	Type: integer
Description: The status of the fan.	
Expected values: <i>ok</i> (i.e. 3)	

Monitoring the System

SOLIDserver relies on common SNMP embedded MIB extensions. Many plug-ins on GitHub, developed for Centreon, are Nagios compatible and can be used with SOLIDserver¹:

- **CPU load:** *loadaverage.pm*
- **Memory usage:** *memory.pm*
- **Swap usage:** *swap.pm*
- **Disk usage:** *check_disk_snmp.pl* - Note that *diskusage.pm* plug-in is not compatible with SOLIDserver embedded MIB.
- **Network traffic:** *interfaces.pm*
- **Running processes:** *processcount.pm*

These metrics should be graphed over time for troubleshooting purposes.

CPU Core Usage

All the OIDs described in this section are part of the MIB extension **HOST-RESOURCES-MIB**.

Table I.11. CPU Core Usage Average 1 min

.1.3.6.1.2.1.25.3.3.1.2.*	
Variable name: hrProcessorLoad	Type: integer
Description: The average percentage of time the CPU core was not idle over the last minute.	

CPU(s) Load

All the OIDs described in this section are part of the MIB extension **UCD-SNMP-MIB**.

Table I.12. CPU(s) Load Average 5 min

.1.3.6.1.4.1.2021.10.1.6.2.*	
Variable name: laLoadInt	Type: integer
Description: The load of all the system's CPUs on an average of 5 min.	
Expected values: Value should remain below the number of CPU cores	
.1.3.6.1.4.1.2021.10.1.6.2	
Variable name: CPUcount	Type: integer
Description: The number of CPU cores is equal to the number of entries returned by an <i>snmpwalk</i> command on this OID.	

Memory Usage

All the OIDs described in this section are part of the MIB extension **UCD-SNMP-MIB**.

¹<https://github.com/centreon/centreon-plugins>

Table I.13. Memory usage

<i>.1.3.6.1.4.1.2021.4.5.0</i>	
Variable name: memTotalReal	Type: integer
Description: The total amount of physical memory in bytes.	
<i>.1.3.6.1.4.1.2021.4.6.0</i>	
Variable name: memAvailReal	Type: integer
Description: The amount of physical available memory in bytes.	
<i>.1.3.6.1.4.1.2021.4.15.0</i>	
Variable name: memCached	Type: integer
Description: The amount of cached memory in bytes.	
<i>.1.3.6.1.4.1.2021.4.14.0</i>	
Variable name: memBuffer	Type: integer
Description: The amount of buffer memory in bytes.	

Swap Usage

All the OIDs described in this section are part of the MIB extension **UCD-SNMP-MIB**.

Table I.14. Swap usage

<i>.1.3.6.1.4.1.2021.4.3.0</i>	
Variable name: memTotalSwap	Type: integer
Description: The total amount of swap space in bytes.	
<i>.1.3.6.1.4.1.2021.4.4.0</i>	
Variable name: memAvailSwap	Type: integer
Description: The amount of swap space currently unused or available in bytes.	
<i>.1.3.6.1.4.1.2021.4.16</i>	
Variable name: memUsedSwapTXT	Type: integer
Description: The amount of swap space currently in use in bytes.	

Disk IO

SOLIDserver is deployed on top of the physical RAID controller. In the *ucdDiskIOMIB*, *diskIOTable* lists all available devices.

The RAID virtual drive should be identified by the following names: *ad[0-9]*, *ada[0-9]*, *mfid[0-9]*.*. Other listed devices should be ignored.

Any sudden and persistent increase in the metrics should be noticed.

All the OIDs described in this section are part of the MIB extension **UCD-SNMP-MIB**.

Table I.15. Disk IO

<i>.1.3.6.1.4.1.2021.13.15.1.1.2.*</i>	
Variable name: diskIODevice	Type: octetstring
Description: The name of the device(s).	
<i>.1.3.6.1.4.1.2021.13.15.1.1.5.*</i>	
Variable name: diskIOReads	Type: integer
Description: The number of read accesses from the device(s) since boot.	

.1.3.6.1.4.1.2021.13.15.1.1.6.*	
Variable name: diskIOWrites	Type: integer
Description: The number of write accesses to the device(s) since boot.	

Disk Usage

Usage should remain under 80% for each mount point's total disk size entry, i.e. the system paths `/`, `/dev`, `/tmp`, `/var`, `/proc` or `/data1`.

Any sudden and rapid increase in the metrics should be noticed.

All the OIDs described in this section are part of the MIB extension **HOST-RESOURCES-MIB**.

Table I.16. Disk usage

.1.3.6.1.2.1.25.2.3.1.3.*	
Variable name: hrStorageDescr	Type: octetstring
Description: A description of the type and instance of the storage described by this entry.	
.1.3.6.1.2.1.25.2.3.1.4.*	
Variable name: hrStorageAllocationUnits	Type: integer
Description: The size of the data objects allocated from this pool in bytes. If this entry is monitoring sectors, blocks, buffers, or packets, for example, this number will commonly be greater than one. Otherwise this number will typically be one.	
.1.3.6.1.2.1.25.2.3.1.5.*	
Variable name: hrStorageSize	Type: integer
Description: The size of the storage represented by this entry, in units of <i>hrStorageAllocationUnits</i> . This object is writable to allow remote configuration of the size of the storage area in the cases where such an operation makes sense and is possible on the underlying system. For example, the amount of main memory allocated to a buffer pool might be modified or the amount of disk space allocated to virtual memory might be modified.	
.1.3.6.1.2.1.25.2.3.1.6.*	
Variable name: hrStorageUsed	Type: integer
Description: The amount of the storage represented by this entry that is allocated, in units of <i>hrStorageAllocationUnits</i> .	

Network Traffic

Any NIC with *ifAdminStatus* set to *up* (i.e. 1) should also have an *ifOperStatus* set to *up* (i.e. 1). Any sudden and rapid increase in the metrics of discard or error packet count should be noticed.

All the OIDs described in this section are part of the MIB extension **IF-MIB**.

Table I.17. Network traffic

.1.3.6.1.2.1.2.2.1.7.*	
Variable name: ifAdminStatus	Type: integer
Description: The NIC admin status, i.e. the desired state of the interface. The <i>testing</i> state (i.e. 3) indicates that no operational packets can be passed.	
Expected values: When the system initializes, all interfaces should be <i>down</i> (i.e. 2) until explicit notice from the system to set it <i>up</i> (i.e. 1).	
.1.3.6.1.2.1.2.2.1.8.*	
Variable name: ifOperStatus	Type: integer

Description: The NIC operational status, i.e. the current operational state of the interface. The <i>testing</i> state (i.e. 3) indicates that no operational packets can be passed.	
Expected values: Should be <i>down</i> (i.e. 2) if <i>ifAdminStatus</i> is <i>down</i> (i.e. 2). If <i>ifAdminStatus</i> is <i>up</i> (i.e. 1), <i>ifOperStatus</i> should be <i>up</i> (i.e. 1) when the interface is ready to transmit or receive network traffic, and <i>dormant</i> (i.e. 5) when the interface is waiting for external actions (such as a serial line waiting for an incoming connection). The NIC remains <i>down</i> (i.e. 2) only if there is a fault that prevents it from being <i>up</i> (i.e. 1): it should be in the <i>notPresent</i> (i.e. 6) state if the interface has missing, usually hardware, components.	
.1.3.6.1.2.1.2.2.1.10.*	
Variable name: ifInOctets	Type: integer
Description: The total traffic received through the interface, including framing characters, in bytes.	
.1.3.6.1.2.1.2.2.1.16.*	
Variable name: ifOutOctets	Type: integer
Description: The total traffic transmitted out of the interface, including framing characters, in bytes.	
.1.3.6.1.2.1.2.2.1.13.*	
Variable name: ifInDiscards	Type: integer
Description: The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.	
.1.3.6.1.2.1.2.2.1.19.*	
Variable name: ifOutDiscards	Type: integer
Description: The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.	
.1.3.6.1.2.1.2.2.1.14.*	
Variable name: ifInErrors	Type: integer
Description: The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.	
.1.3.6.1.2.1.2.2.1.20.*	
Variable name: ifOutErrors	Type: integer
Description: The number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.	

Running Processes

SOLIDserver relies on several processes to operate:

- *IPMServer*, *Postgres* and *HTTPd* should always be running on the monitored SOLIDserver.
- *DHCPd* should be listed in the *hrSWRunName* table when running the DHCP service on the monitored SOLIDserver.
- *Named* should be listed in the *hrSWRunName* table when running the DNS service on the monitored SOLIDserver. Exception is made for the Hybrid DNS engine where *Named* (both authoritative and recursive) can be replaced by *NSD* (authoritative) or *Unbound* (recursive).

All the OIDs described in this section are part of the MIB extension **HOST-RESOURCES-MIB**.

Table I.18. Running processes

.1.3.6.1.2.1.25.4.2.1.2.<PID>	
Variable name: hrSWRunName	Type: octetstring

Description: The process name, i.e. a textual description of the process, including the manufacturer, revision, and the name by which it is commonly known.	
.1.3.6.1.2.1.25.4.2.1.4.<PID>	
Variable name: hrSWRunPath	Type: octetstring
Description: The process binary path, i.e. a description of the location on long-term storage, e.g. a disk drive, from which the process was loaded.	

Monitoring the DHCP Service

SOLIDserver relies on *DHCPd* to provide its DHCP service and metrics on the number of *inform*, *decline*, *release*, *request* and *discover* messages received as well as *ack*, *nack* and *offer* answers sent.

The EfficientIP proprietary MIB EIP-STATS references values as Integers, yet, they represent counters. As in any counter, when the previous value is greater than the current one, this means that the counter has either looped or has been reset. Therefore, it's necessary to interpret these values properly:

- Value > 0 : correct value = Value
- Value < 0 : correct value = $MAXSIGNEDINT^2 + (1 + ABS(MINSIGNEDINT^3) + value)$

These metrics should be graphed over time for troubleshooting purposes.

DHCP service

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.19. DHCP service

.1.3.6.1.4.1.2440.1.3.2.22.1.2.3.97.99.107	
Variable name: ack	Type: integer
Description: The number of ACK answers sent.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.4.110.97.99.107	
Variable name: nack	Type: integer
Description: The number of NACK answers sent.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.5.111.102.102.101.114	
Variable name: offer	Type: integer
Description: The number of OFFER answers sent.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.6.105.110.102.111.114.109	
Variable name: inform	Type: integer
Description: The number of INFORM messages received.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.7.100.101.99.108.105.110.101	
Variable name: decline	Type: integer
Description: The number of DECLINE messages received.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.7.114.101.108.101.97.115.101	
Variable name: release	Type: integer
Description: The number of RELEASE messages received.	

²MAX Signed INT = 2147483647

³MIN Signed INT = -2147483648

.1.3.6.1.4.1.2440.1.3.2.22.1.3.7.114.101.113.117.101.115.116	
Variable name: request	Type: integer
Description: The number of REQUEST messages received.	
.1.3.6.1.4.1.2440.1.3.2.22.1.3.8.100.105.115.99.111.118.101.114	
Variable name: discover	Type: integer
Description: The number of DISCOVER messages received.	

Monitoring the DNS Service

SOLIDserver relies on BIND to provide its DNS service and metrics on the *number of requests over udp* or *tcp*. The metrics also include *IPv4* and *IPv6 requests received by named* or *sent to forwarder/authoritative servers* as well as *recursion*.

The EfficientIP proprietary MIB EIP-STATS references values as Integers, yet, they represent counters. As in any counter, when the previous value is greater than the current one, this means that the counter has either looped or has been reset. Therefore, it's necessary to interpret these values properly:

- Value > 0 : correct value = Value
- Value < 0 : correct value = $MAXSIGNEDINT^4 + (1 + ABS(MINSIGNEDINT^5)) + value$

These metrics should be graphed over time for troubleshooting purposes.

General Name Server Statistics

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.20. General name server statistics

.1.3.6.1.4.1.2440.1.4.2.3.1.3.3.117.100.112	
Variable name: udp	Type: integer
Description: The number of queries received through UDP.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.3.116.99.112	
Variable name: tcp	Type: integer
Description: The number of queries received through TCP.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.9.114.101.113.117.101.115.116.118.52	
Variable name: requestv4	Type: integer
Description: The number of queries received through IPv4.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.9.114.101.113.117.101.115.116.118.54	
Variable name: requestv6	Type: integer
Description: The number of queries received through IPv6.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.11.114.101.115.95.113.117.101.114.121.118.52	
Variable name: res_queryv4	Type: integer
Description: The number of queries sent to external forwarder or authoritative servers using IPv4 for resolution.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.11.114.101.115.95.113.117.101.114.121.118.54	
Variable name: res_queryv6	Type: integer

⁴MAX Signed INT = 2147483647

⁵MIN Signed INT = -2147483648

Description: The number of queries sent to external forwarder or authoritative servers using IPv6 for resolution.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.9.114.101.99.117.114.115.105.111.110	
Variable name: recursion	Type: integer
Description: The number of queries requiring recursion.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.10.114.101.99.117.114.115.101.114.101.106	
Variable name: recurserej	Type: integer
Description: The number of rejected queries requiring recursion.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.9.100.117.112.108.105.99.97.116.101	
Variable name: duplicate	Type: integer
Description: The number of duplicate queries received.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.100.114.111.112.112.101.100	
Variable name: dropped	Type: integer
Description: The number of dropped queries.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.9.114.101.115.95.114.101.116.114.121	
Variable name: res_retry	Type: integer
Description: The number of queries retried on external forwarder or authoritative servers.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.8.114.101.115.112.111.110.115.101	
Variable name: response	Type: integer
Description: The number of responses sent.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.14.114.101.115.95.114.101.115.112.111.110.115.101.118.52	
Variable name: res_responsev4	Type: integer
Description: The number of responses from external forwarder or authoritative servers using IPv4.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.14.114.101.115.95.114.101.115.112.111.110.115.101.118.54	
Variable name: res_responsev6	Type: integer
Description: The number of responses from external forwarder or authoritative servers using IPv6.	

DNS Resolution Queries Statistics

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.21. DNS resolution queries statistics

.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.48	
Variable name: res_queryrtt0	Type: integer
Description: The number of recursive queries with latency < 10ms.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.49	
Variable name: res_queryrtt1	Type: integer
Description: The number of recursive queries with latency >= 10ms & < 100ms.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.50	
Variable name: res_queryrtt2	Type: integer
Description: The number of recursive queries with latency >= 100ms & < 500ms.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.51	
Variable name: res_queryrtt3	Type: integer
Description: The number of recursive queries with latency >= 500ms & < 800ms.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.52	
Variable name: res_queryrtt4	Type: integer

Description: The number of recursive queries with latency \geq 800ms & $<$ 1600ms.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.13.114.101.115.95.113.117.101.114.121.114.116.116.53	
Variable name: res_queryrtt5	Type: integer
Description: The number of recursive queries with latency \geq 1600ms.	

DNS Answers Statistics

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.22. DNS answers statistics

.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.115.117.99.99.101.115.115	
Variable name: success	Type: integer
Description: The number of queries completed successfully, returning the message NOERROR (RCODE: 0).	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.102.111.114.109.101.114.114	
Variable name: formerr	Type: integer
Description: The number of queries with a wrong format, returning the message FORMERR (RCODE: 1).	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.8.115.101.114.118.102.97.105.108	
Variable name: servfail	Type: integer
Description: The number of queries the server failed to complete, returning the message SERVFAIL (RCODE: 2).	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.8.110.120.100.111.109.97.105.110	
Variable name: nxdomain	Type: integer
Description: The number of queries where the domain name does not exist, returning the message NXDOMAIN (RCODE: 3).	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.110.120.114.114.115.101.116	
Variable name: nxrrset	Type: integer
Description: The number of queries for which the RR set should exist but does not, returning the message NXRRSET (RCODE: 8).	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.102.97.105.108.117.114.101	
Variable name: failure	Type: integer
Description: The number of queries failed for which the device sent another RCODE.	

DNS Transfer Requests Statistics

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.23. DNS transfer requests statistics

.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.120.102.114.100.111.110.101	
Variable name: xfrdone	Type: integer
Description: The number of transfer queries successfully completed.	
.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.102.111.114.109.101.114.114	
Variable name: xfrrej	Type: integer
Description: The number of transfer queries rejected.	

DNSSEC Validation Statistics

All the OIDs described in this section are part of the MIB extension **EIP-STATS**.

Table I.24. DNSSEC validation statistics

<i>.1.3.6.1.4.1.2440.1.4.2.3.1.3.7.114.101.115.95.118.97.108</i>	
Variable name: res_val	Type: integer
Description: The total number of DNSSEC validation attempts.	
<i>.1.3.6.1.4.1.2440.1.4.2.3.1.3.14.114.101.115.95.118.97.108.115.117.99.99.101.115.115</i>	
Variable name: res_valsuccess	Type: integer
Description: The total number of DNSSEC validation successfully completed.	
<i>.1.3.6.1.4.1.2440.1.4.2.3.1.3.17.114.101.115.95.118.97.108.110.101.103.115.117.99.99.101.115</i>	
Variable name: res_valnegsuccess	Type: integer
Description: The total number of DNSSEC NX validation successfully completed.	
<i>.1.3.6.1.4.1.2440.1.4.2.3.1.3.11.114.101.115.95.118.97.108.102.97.105.108</i>	
Variable name: res_valfail	Type: integer
Description: The total number of DNSSEC validation failures.	

Monitoring DNS Guardian

DNS Guardian provides advanced security through real time DNS traffic analysis and various SNMP metrics relative to the cache, clients, queries, answers, protection modes, filters and network traffic in general. For more details, refer to the part [Guardian](#).

Each section describes OIDs regarding DNS Guardian cache and the cache of the its views. The only way of monitoring DNS Guardian views is to know their identifier (ID). In the tables below we provide the OIDs until the <view-ID> that you must provide as it depends on your configuration. <view-ID> is an integer between 0 and 7.

Note that you can display additional statistics thanks to:

- The command `help stats`.
- The command `help stats view=<view-ID>`.

DNS Guardian metrics should be graphed over time for troubleshooting purposes.

Cache Size

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.25. Cache size

<i>.1.3.6.1.4.1.2440.1.11.2.4.5.0</i>	
Variable name: eipDNSGUARDIANStatCacheSize	Type: gauge32
Description: The total number of entries in the cache.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.5.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatCacheSize	Type: gauge32
Description: The total number of entries in the cache of the specified view.	

Client Size

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.26. Client size

<i>.1.3.6.1.4.1.2440.1.11.2.4.45.0</i>	
Variable name: eipDNSGUARDIANStatClientsSize	Type: gauge32
Description: The total number of tracked client IP addresses.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.45.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatClientsSize	Type: gauge32
Description: The total number of tracked client IP addresses in the specified view.	

Cache Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.27. Cache statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.3.0</i>	
Variable name: eipDNSGUARDIANStatCacheHit	Type: counter64
Description: The total number of cache hits (including hits in Quarantine and Rescue modes).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.3.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatCacheHit	Type: counter64
Description: The total number of cache hits (including hits in Quarantine and Rescue modes) in the cache of the specified view.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.4.0</i>	
Variable name: eipDNSGUARDIANStatCacheMiss	Type: counter64
Description: The total number of cache misses (including misses in Quarantine and Rescue modes).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.4.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatCacheMiss	Type: counter64
Description: The total number of cache misses (including misses in Quarantine and Rescue modes) in the cache of the specified view.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.10.0</i>	
Variable name: eipDNSGUARDIANStatCacheMissExist	Type: counter64
Description: The total number of queries that did not hit the cache because the related entry has expired.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.10.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatCacheMissExist	Type: counter64
Description: The total number of queries that did not hit the cache of the specified view because the related entry has expired.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.11.0</i>	
Variable name: eipDNSGUARDIANStatCacheMissNotExist	Type: counter64
Description: The total number of queries that did not hit the cache because the related entry doesn't exist.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.11.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatCacheMissNotExist	Type: counter64
Description: The total number of queries that did not hit the cache because the related entry doesn't exist in the cache of the specified view.	

Quarantine Mode Cache Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.28. Quarantine mode cache statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.37.0</i>	
--	--

	Variable name: eipDNSGUARDIANStatCacheHitQuarantine	Type: counter64
	Description: The total number of cache hits in Quarantine mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.37.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheHitQuarantine	Type: counter64
	Description: The total number of cache hits in Quarantine mode in the cache of the specified view.	
.1.3.6.1.4.1.2440.1.11.2.4.34.0		
	Variable name: eipDNSGUARDIANStatCacheMissQuarantine	Type: counter64
	Description: The total number of cache misses in Quarantine mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.34.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheMissQuarantine	Type: counter64
	Description: The total number of cache misses in Quarantine mode in the cache of the specified view.	
.1.3.6.1.4.1.2440.1.11.2.4.35.0		
	Variable name: eipDNSGUARDIANStatCacheMissExistQuarantine	Type: counter64
	Description: The total number of queries that did not hit the cache because the related entry has expired in Quarantine mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.35.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheMissExistQuarantine	Type: counter64
	Description: The total number of queries that did not hit the cache of the specified view because the related entry has expired in Quarantine mode.	
.1.3.6.1.4.1.2440.1.11.2.4.36.0		
	Variable name: eipDNSGUARDIANStatCacheMissNotExistQuarantine	Type: counter64
	Description: The total number of queries that did not hit the cache because the related entry doesn't exist in Quarantine mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.36.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheMissNotExistQuarantine	Type: counter64
	Description: The total number of queries that did not hit the cache of the specified view because the related entry doesn't exist in Quarantine mode.	

Rescue Mode Cache Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.29. Rescue mode cache statistics

.1.3.6.1.4.1.2440.1.11.2.4.43.0		
	Variable name: eipDNSGUARDIANStatCacheHitRescue	Type: counter64
	Description: The total number of cache hits in Rescue mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.43.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheHitRescue	Type: counter64
	Description: The total number of cache hits in Rescue mode in the cache of the specified view.	
.1.3.6.1.4.1.2440.1.11.2.4.40.0		
	Variable name: eipDNSGUARDIANStatCacheMissRescue	Type: counter64
	Description: The total number of cache misses in Rescue mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.40.<view-ID>		
	Variable name: eipDNSGUARDIANViewStatCacheMissRescue	Type: counter64
	Description: The total number of cache misses in Rescue mode in the cache of the specified view.	
.1.3.6.1.4.1.2440.1.11.2.4.41.0		
	Variable name: eipDNSGUARDIANStatCacheMissExistRescue	Type: counter64

Description: The total number of queries that did not hit the cache because the related entry has expired in Rescue mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.41.<view-ID>	
Variable name: eipDNSGUARDIANViewStatCacheMissExistRescue	Type: counter64
Description: The total number of queries that did not hit the cache of the specified view because the related entry has expired in Rescue mode.	
.1.3.6.1.4.1.2440.1.11.2.4.42.0	
Variable name: eipDNSGUARDIANStatCacheMissNotExistRescue	Type: counter64
Description: The total number of queries that did not hit the cache because the related entry doesn't exist in Rescue mode.	
.1.3.6.1.4.1.2440.1.11.2.3.1.42.<view-ID>	
Variable name: eipDNSGUARDIANViewStatCacheMissNotExistRescue	Type: counter64
Description: The total number of queries that did not hit the cache of the specified view because the related entry doesn't exist in Rescue mode.	

Filters Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.30. Filters statistics

.1.3.6.1.4.1.2440.1.11.2.4.47.0	
Variable name: eipDNSGUARDIANStatRatelimitedQuery	Type: counter64
Description: The number of queries per second dropped because source clients have reached the configured query limit rate of 100 queries/s.	
.1.3.6.1.4.1.2440.1.11.2.3.1.47.<view-ID>	
Variable name: eipDNSGUARDIANViewStatRatelimitedQuery	Type: counter64
Description: The number of queries per second dropped because source clients matching the specified view have reached the configured query limit rate of 100 queries/s.	
.1.3.6.1.4.1.2440.1.11.2.4.44.0	
Variable name: eipDNSGUARDIANStatBlockedQuery	Type: counter64
Description: The number of queries per second that were blocked once a trigger configured with the action <i>Block</i> on the related sources has reached its threshold.	
.1.3.6.1.4.1.2440.1.11.2.3.1.44.<view-ID>	
Variable name: eipDNSGUARDIANViewStatBlockedQuery	Type: counter64
Description: The number of queries per second that were blocked once a trigger configured with the action <i>Block</i> on the related sources of the view has reached its threshold.	
.1.3.6.1.4.1.2440.1.11.2.4.2.<trigger-ID>.0	
Variable name: eipDNSGUARDIANStatTrigger<trigger-ID>Armed	Type: counter64
Description: The number of times each specified trigger has been armed. Note that <trigger-ID> identifies each of the 64 available triggers, in the <i>Variable name</i> it is an integer between 0 and 63 and in the OID it is an integer between 00 and 63.	
.1.3.6.1.4.1.2440.1.11.2.3.1.2.<trigger-ID>.<view-ID>	
Variable name: eipDNSGUARDIANViewStatTrigger<trigger-ID>Armed	Type: counter64
Description: The number of times each specified trigger has been armed in the specified view.	

Invalid Queries Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.31. Invalid queries statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.14.0</i>	
Variable name: eipDNSGUARDIANStatRecvInvalidDNSPacket	Type: counter64
Description: The total number of invalid queries that were dropped.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.14.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRecvInvalidDNSPacket	Type: counter64
Description: The total number of invalid queries matching the specified view that were dropped.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.15.0</i>	
Variable name: eipDNSGUARDIANStatRecvInvalidDNSByte	Type: counter64
Description: The total traffic of invalid queries dropped in bytes.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.15.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRecvInvalidDNSByte	Type: counter64
Description: The total traffic of invalid queries matching the specified view dropped in bytes.	

DNS Traffic Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.32. DNS traffic statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.8.0</i>	
Variable name: eipDNSGUARDIANStatRecvDNSPacket	Type: counter64
Description: The total number of received queries.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.8.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRecvDNSPacket	Type: counter64
Description: The total number of received queries matching the specified view.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.6.0</i>	
Variable name: eipDNSGUARDIANStatSendDNSPacket	Type: counter64
Description: The total number of sent queries.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.6.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatSendDNSPacket	Type: counter64
Description: The total number of sent queries matching the specified view.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.9.0</i>	
Variable name: eipDNSGUARDIANStatRecvDNSByte	Type: counter64
Description: The total incoming traffic in bytes.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.9.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRecvDNSByte	Type: counter64
Description: The total incoming traffic matching the specified view in bytes.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.7.0</i>	
Variable name: eipDNSGUARDIANStatSendDNSByte	Type: counter64
Description: The total outgoing traffic in bytes.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.7.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatSendDNSByte	Type: counter64
Description: The total outgoing traffic matching the specified view in bytes.	

DNS Resolution Queries Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.33. DNS resolution queries statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.48.0</i>	
Variable name: eipDNSGUARDIANStatRTT10	Type: counter64
Description: The number of queries with a latency < 10ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.48.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTT10	Type: counter64
Description: The number of queries matching the specified view with a latency < 10ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.49.0</i>	
Variable name: eipDNSGUARDIANStatRTT100	Type: counter64
Description: The number of queries with a latency >= 10ms & < 100ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.49.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTT100	Type: counter64
Description: The number of queries matching the specified view with a latency >= 10ms & < 100ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.50.0</i>	
Variable name: eipDNSGUARDIANStatRTT500	Type: counter64
Description: The number of recursive queries with a latency >= 100ms & < 500ms.	
<i>1.3.6.1.4.1.2440.1.11.2.3.1.50.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTT500	Type: counter64
Description: The number of recursive queries matching the specified view with a latency >= 100ms & < 500ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.51.0</i>	
Variable name: eipDNSGUARDIANStatRTT800	Type: counter64
Description: The number of queries with a latency >= 500ms & < 800m.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.51.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTT800	Type: counter64
Description: The number of queries matching the specified view with a latency >= 500ms & < 800m.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.52.0</i>	
Variable name: eipDNSGUARDIANStatRTT1600	Type: counter64
Description: The number of queries with a latency >= 800ms & < 1600ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.52.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTT1600	Type: counter64
Description: The number of queries matching the specified view with a latency >= 800ms & < 1600ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.4.53.0</i>	
Variable name: eipDNSGUARDIANStatRTTMax	Type: counter64
Description: The number of queries with a latency >= 1600ms.	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.53.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatRTTMax	Type: counter64
Description:	

DNS Answers Statistics

All the OIDs described in this section are part of the MIB extension **EIP-DNSGUARDIAN**.

Table I.34. DNS answers statistics

<i>.1.3.6.1.4.1.2440.1.11.2.4.101.0</i>	
Variable name: eipDNSGUARDIANStatReplyNOERROR	Type: counter64
Description: The number of successful queries, returning the message NOERROR (RCODE: 0).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.101.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNOERROR	Type: counter64
Description: The number of successful queries matching the specified view, returning the message NOERROR (RCODE: 0).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.102.0</i>	
Variable name: eipDNSGUARDIANStatReplyFORMERR	Type: counter64
Description: The number of queries with a wrong format, returning the message FORMERR (RCODE: 1).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.102.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyFORMERR	Type: counter64
Description: The number of queries matching the specified view with a wrong format, returning the message FORMERR (RCODE: 1).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.103.0</i>	
Variable name: eipDNSGUARDIANStatReplySERVFAIL	Type: counter64
Description: The number of queries the server failed to complete, returning the message SERVFAIL (RCODE: 2).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.103.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplySERVFAIL	Type: counter64
Description: The number of queries matching the specified view the server failed to complete, returning the message SERVFAIL (RCODE: 2).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.104.0</i>	
Variable name: eipDNSGUARDIANStatReplyNXDOMAIN	Type: counter64
Description: The number of queries where the domain name does not exist, returning the message NXDOMAIN (RCODE: 3).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.104.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNXDOMAIN	Type: counter64
Description: The number of queries matching the specified view where the domain name does not exist, returning the message NXDOMAIN (RCODE: 3).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.105.0</i>	
Variable name: eipDNSGUARDIANStatReplyNOTIMP	Type: counter64
Description: The number of queries for which the server did not implement the function, returning the message NOTIMP (RCODE: 4).	
<i>1.3.6.1.4.1.2440.1.11.2.3.1.105.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNOTIMP	Type: counter64
Description: The number of queries matching the specified view for which the server did not implement the function, returning the message NOTIMP (RCODE: 4).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.106.0</i>	
Variable name: eipDNSGUARDIANStatReplyREFUSED	Type: counter64
Description: The number of queries the server refused to answer, returning the message REFUSED (RCODE: 5).	
<i>1.3.6.1.4.1.2440.1.11.2.3.1.106.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyREFUSED	Type: counter64
Description: The number of queries matching the specified view the server refused to answer, returning the message REFUSED (RCODE: 5).	

SNMP Metrics

<i>.1.3.6.1.4.1.2440.1.11.2.4.107.0</i>	
Variable name: eipDNSGUARDIANStatReplyYXDOMAIN	Type: counter64
Description: The number of queries for which the name exists when it should not, returning the message YXDOMAIN (RCODE: 6).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.107.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyYXDOMAIN	Type: counter64
Description: The number of queries matching the specified view for which the name exists when it should not, returning the message YXDOMAIN (RCODE: 6).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.108.0</i>	
Variable name: eipDNSGUARDIANStatReplyYXRRSET	Type: counter64
Description: The number of queries for which the RR set exists when it should not, returning the message YXRRSET (RCODE: 7).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.108.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyYXRRSET	Type: counter64
Description: The number of queries matching the specified view for which the RR set exists when it should not, returning the message YXRRSET (RCODE: 7).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.109.0</i>	
Variable name: eipDNSGUARDIANStatReplyNXRRSET	Type: counter64
Description: The number of queries for which the RR set should exist but does not, returning the message NXRRSET (RCODE: 8).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.109.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNXRRSET	Type: counter64
Description: The number of queries matching the specified view for which the RR set should exist but does not, returning the message NXRRSET (RCODE: 8).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.120.0</i>	
Variable name: eipDNSGUARDIANStatReplyNOTAUTH	Type: counter64
Description: The number of queries for which the server is not authoritative for the zone, returning the message NOTAUTH (RCODE: 9).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.120.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNOTAUTH	Type: counter64
Description: The number of queries matching the specified view for which the server is not authoritative for the zone, returning the message NOTAUTH (RCODE: 9).	
<i>.1.3.6.1.4.1.2440.1.11.2.4.121.0</i>	
Variable name: eipDNSGUARDIANStatReplyNOTZONE	Type: counter64
Description: The number of queries for which the name is not in the zone, returning the message NOTZONE (RCODE: 10).	
<i>.1.3.6.1.4.1.2440.1.11.2.3.1.121.<view-ID></i>	
Variable name: eipDNSGUARDIANViewStatReplyNOTZONE	Type: counter64
Description: The number of queries matching the specified view for which the name is not in the zone, returning the message NOTZONE (RCODE: 10).	

Appendix J. Class Studio Pre-defined Variables

This appendix provides a list of the available Class Studio pre-defined variables. For more details regarding the addition details, refer to the chapter Class Studio, in the section [Pre-defined variable](#).

Table J.1. Predefined Variables Classes

Name	The variable must be used in
USER_SOURCE_TYPE	Rights & delegation user classes.
USER_HIDE_PARAM	Rights & delegation user classes.
IP_MANDATORY_MAC_ADDR	IPAM address and address (v6) classes.
IP_NOT_EDITABLE_MAC_ADDR	IPAM address and address (v6) classes.
DHCP_STATIC_NOT_EDITABLE_MAC_ADDR	DHCP static and DHCPv6 static classes.
WORKFLOW_REQUEST_HIDE_ACTION	Workflow request classes.
WORKFLOW_REQUEST_HIDE_ATTACH_TO	Workflow request classes.
WORKFLOW_ADD_TICKET_SPACE	IPAM space classes.
WORKFLOW_ADD_TICKET_BLOCK	IPAM block-type network classes, IPv4 only.
WORKFLOW_ADD_TICKET_SUBNET	IPAM subnet-type network classes, IPv4 only.
WORKFLOW_ADD_TICKET_POOL	IPAM pool classes, IPv4 only.
WORKFLOW_ADD_TICKET_ADDRESS	IPAM address classes, IPv4 only.
WORKFLOW_ADD_TICKET_DNSZONE	DNS zones classes.
FORCE_SUBNET_PREFIX	IPAM subnet-type network classes, in IPv4 and IPv6.
HIDE_IP_ALIAS	IPAM address classes, in IPv4 and IPv6.
HOSTDEV_IS_SWITCH	Device Manage device classes.
NO_SPACE_FATHER_VLSM	IPAM space classes.
NO_VLSM_SUBNET	IPAM subnet-type network classes, in IPv4 and IPv6.
BLOCK_TYPE	IPAM block-type network classes, IPv4 only.

To properly configure the pre-defined variables value and understand the purpose of each one of them, follow the description below.

USER_SOURCE_TYPE

This variable allows to specify the user source.

Value: *local, param, pam or rule.*

USER_HIDE_PARAM

This variable allows to create users using only a login (through the field *usr login*). On the user addition wizard, the *Password, Confirm password, Email, Login URL* and field *Maintainer groups* are hidden.

Value: *1* (one) to enable the variable. Leave the field empty to disable it.

IP_MANDATORY_MAC_ADDR

This variable allows to make the MAC address field required in the IPv4 and IPv6 address addition and edition wizards.

Value: *1* (one) to enable the variable. Leave the field empty to disable it.

IP_NOT_EDITABLE_MAC_ADDR

This variable allows to prevent users from editing the MAC address field in the IPv4 and IPv6 address addition and edition wizards.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

DHCP_STATIC_NOT_EDITABLE_MAC_ADDR

This variable allows to prevent users from editing the MAC address field in the DHCP static addition and edition wizards.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_REQUEST_HIDE_ACTION

This variable allows to hide the *Action requested* field in the Workflow outgoing requests addition wizard.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_REQUEST_HIDE_ATTACH_TO

This variable allows to hide the *Attach to* drop-down list in the Workflow outgoing requests addition wizard.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_SPACE

This variable allows to associate a Workflow request with a space directly from the addition/edition wizard in the IPAM.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_BLOCK

This variable allows to associate a Workflow request with a block-type network directly from the addition/edition wizard in the IPAM.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_SUBNET

This variable allows to associate a Workflow ticket with a subnet-type network directly from the addition/edition wizard in the IPAM.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_POOL

This variable allows to associate a Workflow ticket with a pool directly from the addition/edition wizard in the IPAM.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_ADDRESS

This variable allows to associate a Workflow ticket with an IP address directly from the addition/edition wizard in the IPAM.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

WORKFLOW_ADD_TICKET_DNSZONE

This variable allows to associate a Workflow ticket with a zone directly from the addition/edition wizard in the DNS.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

FORCE_SUBNET_PREFIX

This variable allows to force the value of a terminal network prefix in the addition wizard. For more details, refer to the section [Force prefix](#).

Value: the prefix of your choice following the format *<number>*. Leave the field empty to disable the variable.

VLSM Specificity: if you set this variable on a non-terminal subnet-type network, it only applies to the non-terminal subnet-type network itself, it does not apply to the networks it contains.

HIDE_IP_ALIAS

This variable allows to hide the *Aliases configuration* page in the IP address addition wizard. For more details, refer to the section [Configuring and Managing IP Address Aliases](#).

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

HOSTDEV_IS_SWITCH

This variable allows to specify that a device is a switch in the module Device manager.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

NO_SPACE_FATHER_VLSM

This variable allows to prevent a space from being affiliated with a parent space in the space addition wizard.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

NO_VLSM_SUBNET

This variable allows to prevent users from setting a terminal network as non-terminal in the network addition wizard. The box *Terminal network* is hidden.

Value: 1 (one) to enable the variable. Leave the field empty to disable it.

BLOCK_TYPE

This variable allows to set manually the start and end address of a block-type network, like you are able to when creating DHCP ranges or IPAM pools.

Value: *range* to enable the variable. Leave the field empty to disable it.

Appendix K. Configuring RADIUS

Table of Contents

Configuring FreeRADIUS	1289
Configuring RADIUS with Cisco ACS	1290
Configuring OneTime Password with Token Authentication	1291

Configuring FreeRADIUS

If you intend to authenticate users via RADIUS, you can configure FreeRadius to retrieve your groups of users. Once FreeRadius is configured, do not forget to add the RADIUS users authentication rule as detailed in the section [Relying on RADIUS Authentication](#).

Configuring the RADIUS Server

To start with, your RADIUS server must be configured with the following information:

- The addresses of the SOLIDserver appliance (the RADIUS 'clients') that might connect to it.
- The EfficientIP vendor number: 2440.
- The dictionary *efficientip*, it must be configured to send back the following attributes:

Table K.1. Efficientip dictionary: attributes to be returned

Attribute	Code	Type
efficientip-version	1	integer
efficientip-service-class	2	integer
efficientip-identity-type	3	integer
efficientip-first-name	16	string
efficientip-last-name	17	string
efficientip-pseudonym	18	string
efficientip-ip-host	19	string
efficientip-email	20	string
efficientip-first-login-path	32	string
efficientip-maintainer-group	33	string
efficientip-groups	34	string
efficientip-admin-group	35	string
efficientip-extra-blob	64	string

Configuring a FreeRadius server with SOLIDserver

To send group of users to SOLIDserver and complete the configuration, your FreeRadius server needs three files: *dictionary.efficientip*, *clients.conf* and *users*.

These files should at least contain the following information:

dictionary.efficientip

```
#Dictionary for efficientip
VENDOR efficientip 2440

BEGIN-VENDOR efficientip

ATTRIBUTE efficientip-version 1 integer
ATTRIBUTE efficientip-service-class 2 integer
ATTRIBUTE efficientip-identity-type 3 integer
ATTRIBUTE efficientip-first-name 16 string
ATTRIBUTE efficientip-last-name 17 string
ATTRIBUTE efficientip-pseudonym 18 string
ATTRIBUTE efficientip-ip-host 19 string
ATTRIBUTE efficientip-email 20 string
ATTRIBUTE efficientip-first-login-path 32 string
ATTRIBUTE efficientip-maintainer-group 33 string
ATTRIBUTE efficientip-groups 34 string
ATTRIBUTE efficientip-admin-group 35 string
ATTRIBUTE efficientip-extra-blob 64 string

END-VENDOR efficientip
```

clients.conf

```
client SDS-1000 {
    ipaddr = 192.168.100.100
    secret = abc123
}
```

users

```
localuser Cleartext-Password := "Password123"
efficientip-groups = "mygroup",
```

Configuring RADIUS with Cisco ACS

You can configure the RADIUS authentication rule to authenticate users against a Cisco Secure Access Control Server (ACS). To do so, define the EfficientIP RADIUS vendor and VSA set in a RADIUS vendor/VSA import file named `efficientip.ini`, then follow the procedure below.

To configure Cisco ACS with RADIUS

1. On the computer running ACS, open an MS-DOS command prompt.
2. Change directories until you get to the directory containing the file `CSUtil.exe`. For more details regarding its location, refer to the Cisco ACS documentation.
3. Once you are in the right directory, execute the command below:

```
CSUtil.exe -addUDV 5 efficientip.ini
```

In this command, the number 5 is an unused ACS RADIUS vendor slot number and `efficientip.ini` is the name of the EfficientIP RADIUS vendor/VSA import file you created earlier.

4. Press Enter. A `CSUtil.exe` confirmation prompt appears.
5. Confirm that you want to add the RADIUS vendor and halt all ACS services during the process, type in Y and press Enter. `CSUtil.exe` halts ACS services, parses the vendor/VSA input file, and adds the new RADIUS vendor and VSAs to ACS. This process may take a few minutes. After it is complete, `CSUtil.exe` restarts ACS services.

Example of an import file "efficientip.ini" for RADIUS vendor/VSA where *EfficientIP* is set as a vendor and 2440 is the IETF code number:

```
[User Defined Vendor] Name=EfficientIP IETF Code=2440
VSA 1=efficientip-version
VSA 2=efficientip-service-class
VSA 3=efficientip-identity-type
VSA 16=efficientip-first-name
VSA 17=efficientip-last-name
VSA 18=efficientip-pseudonym
VSA 19=efficientip-ip-host
VSA 20=efficientip-email
VSA 32=efficientip-first-login-path
VSA 33=efficientip-maintainer-group
VSA 34=efficientip-groups
VSA 35=efficientip-admin-group
VSA 64=efficientip-extra-blob

[efficientip-version]
Type=INTEGER
Profile=OUT

[efficientip-service-class]
Type=INTEGER
Profile=OUT

[efficientip-identity-type]
Type=INTEGER
Profile=OUT

[efficientip-first-name]
Type=STRING
Profile=OUT

[efficientip-last-name]
Type=STRING
Profile=OUT

[efficientip-pseudonym]
Type=STRING
Profile=OUT

[efficientip-ip-host]
Type=STRING
Profile=OUT

[efficientip-email]
Type=STRING
Profile=OUT

[efficientip-first-login-path]
Type=STRING
Profile=OUT

[efficientip-maintainer-group]
Type=STRING
Profile=OUT

[efficientip-groups]
Type=STRING
Profile=MULTI OUT

[efficientip-admin-group]
Type=STRING
Profile=OUT

[efficientip-extra-blob]
Type=STRING
Profile=OUT
```

Configuring OneTime Password with Token Authentication

You can set up OneTime Password (OTP) based on a token authentication to secure user sessions. OTP associates a user session with a unique password valid for a limited period of time.

SOLIDserver caches the credentials to authenticate every user operation. When the cache expires, SOLIDserver uses the cached credentials to generate and send a new authentication request to the RADIUS server. If this request fails, the user is disconnected.

To control when the client is disconnected, a set of registry database entries allow to define for how long SOLIDserver should cache the data:

1. You can either define for how long SOLIDserver should cache passwords, as detailed in the section [Caching OTP Credentials For a Certain Time](#).
2. Or you can make sure that the cache does not expire while the user is active, as detailed in the section [Renewing Cached OTP Credentials for Logged Users](#).

Prerequisites

- Configuring RADIUS authentication rule. For more details, refer to the section [Relying on RADIUS Authentication](#).
- Belonging to a group *admin*, the only group that can access to the page Registry database.

Caching OTP Credentials For a Certain Time

Once you met the [prerequisites](#), you can configure a registry database key to define for how long SOLIDserver should cache your OTP credentials. To ensure that the user is not disconnected while SOLIDserver is still open, we recommend that you also configure the key that determines the duration of SOLIDserver user sessions.

If you would rather edit one key and ensure your users are not disconnected while their session is active, refer to the section [Renewing Cached OTP Credentials for Logged Users](#).

To define for how long passwords should be cached

Only users of the group *admin* can perform this operation.

1. **Adding the registry key that defines for how long OTP credentials are cached**
 - a. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
 - b. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
 - c. In the menu, click on **+ Add**. The wizard **Registry database Add an item** opens.
 - d. In the field **Name**, type in *ipmserver.login.password_cache_time* .
 - e. In the field **Value**, type in the value of your choice, in seconds. The recommended value is *900*.
 - f. Click on to complete the operation. The report opens and closes. The page refreshes and the new key is listed.
2. **Editing the key that sets SOLIDserver user session time**
 - a. In the search engine of the column **Name**, type in *www.login.session_timeout*. Only this key is listed.
 - b. In the column **Value**, click on the value listed. The wizard **Registry database Edit a value** opens.

- c. In the field **Value**, type in the value of your choice, in seconds. It should be shorter than the value set for the key `ipmserver.login.password_cache_time` to ensure the user OTP credentials do not expire before SOLIDserver session ends.
- d. Click on to complete the operation. The report opens and closes. The page refreshes and the new value is displayed.

Renewing Cached OTP Credentials for Logged Users

Once you met the [prerequisites](#), you can add the registry database key that ensures that logged users authenticated using OTP are not disconnected while their SOLIDserver session is running. This key renews cached credentials as long as the user is active.

To edit the registry key that renews cached credentials while the session is active

Only users of the group `admin` can perform this operation.

1. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
2. In the section **Expert**, click on **Registry database**. The page **Registry database** opens.
3. In the menu, click on  **Add**. The wizard **Registry database Add an item** opens.
4. In the field **Name**, type in `ipmserver.login.password_cache_time_renew` .
5. In the field **Value**, type in 1 to enable the protection of OTP authenticated .
6. Click on to complete the operation. The report opens and closes. The page refreshes and the new key is listed.

Appendix L. Using Remote Authentication for SSH Connections to SOLIDserver

Table of Contents

Configuring LDAP Authentication for SSH Connections	1294
Configuring RADIUS Authentication for SSH Connections	1299

EfficientIP allows to configure SOLIDserver to use LDAP authentication for Secure Shell connections. This configuration allows to grant existing LDAP users access to as many SOLIDserver appliances as you want.

Configuring LDAP Authentication for SSH Connections

To configure the LDAP authentication for SSH connections to SOLIDserver, you must:

1. Understand the [Specificities](#) and [Prerequisites](#).
2. Edit the LDAP configuration file. For more details, refer to the section [Editing the LDAP Configuration File](#).
3. Edit the Name Service Switch configuration file. For more details, refer to the section [Editing the Name Service Switch Configuration File](#).
4. Edit the PAM configuration for SSH daemon. For more details, refer to the section [Editing the PAM Configuration for SSH Daemon](#).
5. If you use an SSL certificate self-signed by Windows, you must also edit the LDAP communication settings. For more details, refer to the section [Editing the LDAP Communication Settings](#).

Prerequisites

- An LDAP server properly configured and running.
- The LDAP server and SOLIDserver must be set at the same time.
 - Make sure the LDAP server is on time.
 - Configure NTP servers on SOLIDserver. If you want LDAP authentication for several SOLIDserver appliances, the NTP must be configured on every appliance. For more details, refer to the section [Configuring NTP Servers](#).
- To set up the LDAP authentication for several appliances, you must:
 1. Configure the authentication from the managing SOLIDserver.
 2. Once the configuration is complete, apply it the remote appliance(s).

Specificities

The configuration of LDAP authentication for SSH connections:

- Must be done from the CLI.
- Must be done locally from a SOLIDserver management appliance. When the configuration is complete:
 - You can apply it to the remote SOLIDserver appliance(s) you manage.
 - It automatically applies to the Hot Standby, if your SOLIDserver is configured with High Availability.
- Allows to grant access to SOLIDserver to existing LDAP users, there is no need to edit SOLIDserver local user or group of users database. That way, if your LDAP server is not responding, local users with sufficient rights can still access the appliance via SSH.

Editing the LDAP Configuration File

You must first edit the LDAP configuration file to specify the class of users you want to grant SSH access to. During that edition you must also define the type of permissions you grant to the users of the class specified.

To edit the file `ldap.conf`

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account *admin*.
3. Edit the file `/data1/etc/ldap.conf` using the following command:

```
% emacs /data1/etc/ldap.conf
```

The edited file should include the following:

```
# Specify the FQDN or IP address of the LDAP server.
host <IP-address-or-hostname>
# Set the base search that suits your needs: cn=BaseSearch, dc=EXAMPLE, dc=COM.
base <your-base-search>
# Set an existing user among the following: cn=admin, cn=BaseSearch, dc=EXAMPLE, dc=COM.
# This parameter cannot be empty otherwise you will not be able to list users.
rootbinddn <user>

# Enable LDAPS if need be. Valid values: on, off.
ssl <status>

# Specify for the following attributes the name of the LDAP user class that can access SOLIDserver.
pam_filter objectclass=<userclass>
nss_map_objectclass posixAccount <userclass>
# Specify for the following attributes the name of the attribute of the selected user class
# that contains the user login.
pam_login_attribute <login-attribute>
nss_map_attribute uid <login-attribute>

# If the specified user class is not already set with the attributes: uidNumber, gidNumber,
# loginShell, homeDirectory.
# You must specify for the following attributes the name of the attribute of the selected user
# class that contains the users uidNumber and gidNumber.
nss_map_attribute uidNumber <UnixUID-attribute>
nss_map_attribute gidNumber <UnixGID-attribute>
# You can specify for the following attributes the name of the attribute of the selected user
# that contains the loginShell and homeDirectory used to connect to your appliance.
# These attributes are optional, so if you do not want to set them either comment the lines or
# do not include them at all. These attributes cannot be declared without value.
nss_map_attribute loginShell <attribute-containing-the-shell-you-connect-to>
nss_map_attribute homeDirectory <attribute-containing-the-path-to-the-directory-"home">

# Set the level of permissions of the LDAP users accessing SOLIDserver via SSH.
# If you set the value of the example below, you grant administrative rights to all the users
# belonging to the classes specified with "pam_filter" and "nss_map_objectclass".
nss_override_attribute_value uidNumber 1001
nss_override_attribute_value gidNumber 1000
```

```
nss_override_attribute_value homeDirectory /data1/users/admin  
nss_override_attribute_value loginShell /bin/csh
```

4. Edit the file `/data1/etc/ldap.secret` to insert your own LDAP password for the account `root-binddn` using the following command:

```
% emacs /data1/etc/ldap.secret
```

Editing the Name Service Switch Configuration File

Once you edited the file `ldap.conf`, you must edit the name service switch configuration file to edit where the authentication information must be retrieved.

To edit the file `nsswitch.conf`

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account `admin`.
3. Edit the file `/data1/etc/nsswitch.conf` using the following command:

```
% emacs /data1/etc/nsswitch.conf
```

4. The edited file should include the following:

```
#  
# nsswitch.conf(5) - name service switch configuration file  
# $FreeBSD: releng/10.2/etc/nsswitch.conf 224765 2011-08-10 20:52:02Z dougb $  
#  
# Edit the value of "group" to retrieve LDAP data. Initial value: "group: compat"  
group: files ldap  
# Comment the line "group_compat: nis"  
# group_compat: nis  
  
hosts: files dns  
networks: files  
  
# Edit the value of "passwd" to retrieve LDAP data. Initial value: "passwd: compat"  
passwd: files ldap  
# Comment the line "passwd_compat: nis"  
# passwd_compat: nis  
  
shells: files  
services: compat  
services_compat: nis  
protocols: files  
rpc: files
```

Editing the PAM Configuration for SSH Daemon

Once you edited the files `ldap.conf` and `nsswitch.conf`, you must edit the PAM configuration of SSH daemon to include LDAP settings.

To edit the file `pam.d/sshd`

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account `admin`.
3. Edit the file `/data1/etc/pam.d/sshd` using the following command:

```
% emacs /data1/etc/pam.d/sshd
```

4. The edited file should include the following:

Using Remote Authentication for SSH Connections to SOLIDserver

```
# $FreeBSD: releng/10.2/etc/pam.d/sshd 197769 2009-10-05 09:28:54Z des $
#
# PAM configuration for the "sshd" service
#
# auth
auth          sufficient      pam_opie.so           no_warn no_fake_prompts
auth          requisite       pam_opieaccess.so    no_warn allow_local
#auth         sufficient      pam_krb5.so          no_warn try_first_pass
#auth         sufficient      pam_ssh.so           no_warn try_first_pass
# Add the following line to specify that LDAP is used for authentication.
auth         sufficient      /usr/local/lib/pam_ldap.so    no_warn
auth          sufficient      pam_unix.so          no_warn try_first_pass

# account
account       required        pam_nologin.so
#account      required        pam_krb5.so
account       required        pam_login_access.so
# Add the following line to specify that you want the credentials to be verified via LDAP.
account      sufficient      /usr/local/lib/pam_ldap.so    no_warn ignore_authinfo_unavail
ignore_unknown_user
account       required        pam_unix.so

# session
#session      optional        pam_ssh.so           want_agent
session       required        pam_permit.so

# password
#password     sufficient      pam_krb5.so          no_warn try_first_pass
password      required        pam_unix.so          no_warn try_first_pass
```

Editing the LDAP Communication Settings

If you use a self-signed certificate generated on MS Windows, you have to edit the LDAP communication settings: the file *openldap.conf* that defines the options to communicate with LDAP.

Once you edited the files *ldap.conf*, *nsswitch.conf* and *pam.d/sshd*, follow the procedure below.

To edit the file *openldap.conf*

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account *admin*.
3. Edit the file */data1/etc/openldap.conf* using the following command:

```
% emacs /data1/etc/openldap.conf
```

4. The edited file should include the following:

```
#
# LDAP Defaults
#
# See ldap.conf(5) for details
# This file should be world readable but not world writable.

#BASE      dc=example,dc=com
#URI       ldap://ldap.example.com ldap://ldap-master.example.com:666

#SIZELIMIT 12
#TIMELIMIT 15
#DEREF     never

# Add the following line to disable the CA check.
TLS_REQCERT never

#Dictionary for efficientip
VENDOR    efficientip 2440
```

Using Remote Authentication for SSH Connections to SOLIDserver

```
BEGIN-VENDOR efficientip
ATTRIBUTE efficientip-version 1 integer
ATTRIBUTE efficientip-service-class 2 integer
ATTRIBUTE efficientip-identity-type 3 integer
ATTRIBUTE efficientip-first-name 16 string
ATTRIBUTE efficientip-last-name 17 string
ATTRIBUTE efficientip-pseudonym 18 string
ATTRIBUTE efficientip-ip-host 19 string
ATTRIBUTE efficientip-email 20 string
ATTRIBUTE efficientip-first-login-path 32 string
ATTRIBUTE efficientip-maintainer-group 33 string
ATTRIBUTE efficientip-groups 34 string
ATTRIBUTE efficientip-admin-group 35 string
ATTRIBUTE efficientip-extra-blob 64 string
END-VENDOR efficientip
```

Once you edited the the LDAP communication settings, you need to [make sure the configuration is properly set](#).

Making Sure the Configuration is Properly Set

Once your configuration is complete, there are two ways of making sure it is properly set:

1. Use the command *getent* to list all the users with SSH access and make sure your LDAP users are listed and that you did not grant access to the unwanted users.
2. Try to connect to SOLIDserver via SSH using the credentials of an LDAP user to which you granted access. They were listed in the command result.

To check the list of users allowed to connect via SSH

1. Open a shell session.
2. Connect to SOLIDserver using the credentials of the account *admin*.
3. Use the following command:

```
% getent passwd
```

This command returns the list of all the users that can connect via SSH. All local, LDAP and/or RADIUS users are listed as follows:

```
<username>:*:1001:1000:admin:/data1/users/admin:/bin/csh
```

Once you edited the made sure the configuration is properly set, you need to [apply the configuration to remote appliances](#).

Applying the Configuration to Remote Appliances

If you manage several SOLIDserver appliances from the page *Centralized Management*, once you configured the LDAP authentication for SSH connections on the management appliance you can push your configuration, i.e. apply it, to remote appliances.

Before applying the configuration, refer to the section [Making Sure the Configuration is Properly Set](#) to ensure you are not pushing an erroneous configuration on your network. If the configuration is incorrect, pushing it might prevent you from connecting to the remote appliance(s) via SSH.

To apply the local LDAP authentication configuration for SSH to remote appliances

1. Open a browser.
2. Connect to the SOLIDserver management appliance you configured using its IP address or hostname.
3. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
4. In the section **System**, click on the button **Centralized Management**. The page **Centralized Management** opens.
5. Tick the remote appliance(s) to which you want to apply the local LDAP authentication.
6. In the menu, select . **Tools > Push local LDAP/RADIUS configuration**. The wizard **Push the local LDAP/RADIUS authentication** opens.
7. Click on to complete the operation. The report opens and closes. The page **Centralized Management** is visible again.

Configuring RADIUS Authentication for SSH Connections

To configure the RADIUS authentication for SSH connections to SOLIDserver, you must:

1. Follow the procedures described in the appendix [Configuring RADIUS](#).
2. Understand the [Specificities](#) and [Prerequisites](#).
3. Edit the RADIUS configuration file. For more details, refer to the section [Editing the LDAP Configuration File](#).
4. Edit the PAM configuration for SSH daemon. For more details, refer to the section [Editing the RADIUS Configuration for SSH Daemon](#).
5. Edit the list of RADIUS users. For more details, refer to the section [Editing the RADIUS Users List](#).

Prerequisites

- A RADIUS server properly configured and running.
- The RADIUS server and SOLIDserver must be set at the same time.
- To set up the RADIUS authentication for several appliances, you must:
 1. Configure the authentication from the managing SOLIDserver.
 2. Once the configuration is complete, apply it the remote appliance(s).
- The user must exist on SOLIDserver.

Specificities

The configuration of RADIUS authentication for SSH connections:

- Must be done from the CLI.
- Must be done locally from a SOLIDserver management appliance. When the configuration is complete:
 - You can apply it to the remote SOLIDserver appliance(s) you manage.
 - It automatically applies to the Hot Standby, if your SOLIDserver is configured with High Availability.

- If your RADIUS server is not responding, local users with sufficient rights can still access the appliance via SSH.

Editing the RADIUS Configuration File

You must first edit the RADIUS configuration file¹ to specify the class of users you want to grant SSH access to. During that edition you must also define the type of permissions you grant to the users of the class specified.

To edit the file `radius.conf`

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account *admin*.
3. Edit the file `/data1/etc/radius.conf` using the following command:

```
% emacs /data1/etc/radius.conf
```

The edited file should include the RADIUS server IP address, secret key, timeout in seconds and the maximum number of attempts, as follows: `<radius_server_ip>:<port> <secret_key> <timeout> <maximum_attempts>`. Not indicating the port, as in the example below, automatically sets it to the default RADIUS port:

```
#radius_server_ip:port secret_key timeout maximum_attempts  
1.2.3.4 RadiusSecretKey 3
```

Editing the RADIUS Configuration for SSH Daemon

Once you edited the file `radius.conf`, you must edit the PAM configuration of SSH daemon to include RADIUS settings.

To edit the file `pam.d/sshd`

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account *admin*.
3. Edit the file `/data1/etc/pam.d/sshd` using the following command:

```
% emacs /data1/etc/pam.d/sshd
```

4. The edited file should include the following:

```
#  
# $FreeBSD: releng/10.2/etc/pam.d/sshd 197769 2009-10-05 09:28:54Z des $  
#  
# PAM configuration for the "sshd" service  
#  
# auth  
auth sufficient pam_opie.so no_warn no_fake_prompts  
auth requisite pam_opieaccess.so no_warn allow_local  
#auth sufficient pam_krb5.so no_warn try_first_pass  
#auth sufficient pam_ssh.so no_warn try_first_pass  
# Radius authentication  
auth sufficient pam_radius.so  
auth required pam_unix.so no_warn try_first_pass  
  
# account  
account required pam_nologin.so  
#account required pam_krb5.so  
account required pam_login_access.so
```

¹For more details, refer to the `radius.conf` Linux man page available at <https://linux.die.net/man/5/radius.conf>.

Using Remote Authentication for SSH Connections to SOLIDserver

```
account          required      pam_unix.so
# session
#session         optional     pam_ssh.so      want_agent
session         required     pam_permit.so

# password
#password        sufficient  pam_krb5.so     no_warn try_first_pass
password        required     pam_unix.so     no_warn try_first_pass
```

Editing the RADIUS Users List

Once you edited the files *radius.conf* and *pam.d/sshd*, you need to add the users for which you want to allow RADIUS authentication.

If you want a user to connect in SSH only via RADIUS authentication, make sure that his password is empty. Such user cannot log in using SSH any other way than using RADIUS since, by default, the server does not allow login to accounts with an empty password.

Any user added to the list will be able to connect SOLIDserver using the secret key configured in the section [Editing the RADIUS Configuration for SSH Daemon](#).

To add a RADIUS user

1. Open a shell session.
2. Connect to a SOLIDserver management appliance using the credentials of the account *admin*.
3. Add a user using the following command:

```
% pw useradd -n <username> -u <uid> -m
```

Note that, as a best practice, setting a *uid* above 10000 allows to easily identify "Radius-only users".

Once you edited the RADIUS users list, you need to [make sure the configuration is properly set](#).

Making Sure the Configuration is Properly Set

Once your configuration is complete, there are two ways of making sure it is properly set:

1. Use the command *getent* to list all the users with SSH access and make sure your RADIUS users are listed and that you did not grant access to the unwanted users.
2. Try to connect to SOLIDserver via SSH using the credentials of a RADIUS user to which you granted access. They were listed in the command result.

To check the list of users allowed to connect via SSH

1. Open a shell session.
2. Connect to SOLIDserver using the credentials of the account *admin*.
3. Use the following command:

```
% getent passwd
```

This command returns the list of all the users that can connect via SSH. All local, LDAP and/or RADIUS users are listed as follows:

```
<username>:*:1001:1000:admin:/data1/users/admin:/bin/csh
```

Once you edited the made sure the configuration is properly set, you need to [apply the configuration to remote appliances](#).

Applying the Configuration to Remote Appliances

If you manage several SOLIDserver appliances from the page *Centralized Management*, once you configured the RADIUS authentication for SSH connections on the management appliance you can push your configuration, i.e. apply it, to remote appliances.

Before applying the configuration, refer to the section [Making Sure the Configuration is Properly Set](#) to ensure you are not pushing an erroneous configuration on your network. If the configuration is incorrect, pushing it might prevent you from connecting to the remote appliance(s) via SSH.

To apply the local RADIUS authentication configuration for SSH to remote appliances

1. Open a browser.
2. Connect to the SOLIDserver management appliance you configured using its IP address or hostname.
3. In the sidebar, click on  **Administration** or **Admin Home**. The page **Admin Home** opens.
4. In the section **System**, click on the button **Centralized Management**. The page **Centralized Management** opens.
5. Tick the remote appliance(s) to which you want to apply the local RADIUS authentication.
6. In the menu, select  **Tools** > **Push local LDAP/RADIUS configuration**. The wizard **Push the local LDAP/RADIUS authentication** opens.
7. Click on to complete the operation. The report opens and closes. The page **Centralized Management** is visible again.

Appendix M. Configuring Non-Supported Options

Table of Contents

Prerequisites	1304
Limitations	1304
Configuring Non-Supported Firewall Rules	1305
Configuring Non-Supported Apache Settings	1307
Configuring Non-Supported Unbound Settings	1309
Configuring Non-Supported NSD Settings	1311
Configuring Non-Supported BIND Settings	1313
Configuring Non-Supported SNMP Settings	1317
Configuring Non-Supported DHCP Configurations	1318
Configuring Non-Supported NTP Settings	1320
Configuring Non-Supported syslog-ng Settings	1321
Configuring Non-Supported PostgreSQL Settings	1322

Administrators can incorporate options - configurations, settings - that are not supported by SOLIDserver via CLI.

Before configuring any non-supported configuration, keep in mind that:

- These configurations must be set **at your own risk**. No support help can be expected after setting any of the non-supported configuration described in this appendix.
- These configurations are **very advanced and must be carefully implemented**.

You can set non-supported configurations for: firewall rules and options for the services Apache, Unbound, NSD, BIND, SNMP, DHCP, NTP, syslog-ng and PostgreSQL.

Prerequisites

- SOLIDserver in version 6.0.1 or greater.
- Configuring local servers: DHCP EfficientIP servers, DNS EfficientIP servers or Hybrid DNS servers.
- The service you want to configure must be running.
- The user configuring the non-supported options must have:
 - Administrative rights over SSH connections to SOLIDserver.
 - A good understanding of the environment and of the services configuration file, syntax and options.
- Checking the changes before applying them to the production environment.

Limitations

- All changes must be performed via SSH, you cannot configure or display non-supported options from the GUI.
- The configuration of non-supported options can only be done in a specific sections of the files and nowhere else.
- You can only configure locally non-supported options.
- You can only configure non-supported options on physical servers EfficientIP DNS or DHCP servers, you cannot set them on a smart architecture.

If you want to add non-supported options on several physical servers managed via the same architecture, you must set them on each server configuration file individually.

- The non-supported options that you configure overwrite the current configuration. So make sure that the options you incorporate are not already set in the configuration because the GUI might not reflect these changes. Besides, configuring options twice may prevent the service from running properly.

Configuring Non-Supported Firewall Rules

You can configure non-supported firewall rules that the `ipfw` files automatically incorporate:

- In restricted mode, non-supported rules can be incorporated in the file `/ipfw/ipfw.rules`.
- In open mode, non-supported rules can be incorporated in the `/ipfw/ipfw_open.rules`.

By default, both files exist in the directory `/ipfw`. Note that the current firewall mode of your appliance, *Restricted* and *Open*, is visible in the GUI on the page *Network configuration*, for more details refer to the section [Setting the Firewall](#).

Before configuring non-supported rules, keep in mind that there is no validation check for these files, so **if you misconfigure either you can lose access to your appliance**, via SSH or otherwise.

Configuring Non-Supported Firewall Configurations in Restricted Mode

By default, the directory `/ipfw` contains a file `ipfw.rules` that can be edited to configure the non-supported firewall rules of your choice if you are in restricted mode.

`ipfw.rules`

```
#!/bin/sh

/sbin/ipfw -q flush

if [ -f /etc/ipfw.rules.stats ]; then
  . /etc/ipfw.rules.stats
fi
if [ -f /etc/ipfw.rules.sourcerouting ]; then
  . /etc/ipfw.rules.sourcerouting
fi

/sbin/ipfw -q add 10 allow ip from any to any via lo0
/sbin/ipfw -q add 14 allow carp from any to any
/sbin/ipfw -q add 20 check-state
...
/sbin/ipfw -q add 59999 deny ip from any to any
/sbin/ipfw -q add 61000 allow ip from any to any

## Include file for customizations
## USE WITH CAUTION
incfile=/usr/local/nessy2/etc/ipfw/ipfw.rules.inc
if [ -f $incfile ]; then
  . $incfile
fi
```



`ipfw.rules.inc`

```
/sbin/ipfw -q add 5 deny icmp from any to any
echo "PING requests are blocked"
```

Figure M.1. Example of non-supported firewall rules in restricted mode

Before configuring non-supported rules, keep in mind that there is no validation check for these files, so **if you misconfigure either you can lose access to your appliance**, via SSH or otherwise.

To incorporate non-supported firewall rules in restricted mode

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and *root* credentials.

3. In the directory `/usr/local/nessy2/etc/ipfw/`, edit the file `ipfw.rules.inc` to incorporate the non-supported firewall rules of your choice.

Keep in mind that if you misconfigure this file you can lose access to your appliance.

4. Restart the firewall daemon to take into account your changes with the command:

```
service ipfw restart
```

Configuring Non-Supported Firewall Configurations in Open Mode

By default, the directory `/ipfw` contains a file `ipfw_open.rules` that can be edited to configure the non-supported firewall rules of your choice if you are in open mode.

ipfw_open.rules

```
#!/bin/sh
/sbin/ipfw -q flush

if [ -f /etc/ipfw.rules.stats ]; then
    . /etc/ipfw.rules.stats
fi
if [ -f /etc/ipfw.rules.sourcerouting ]; then
    . /etc/ipfw.rules.sourcerouting
fi

/sbin/ipfw -q add 00100 allow ip from any to any
/sbin/ipfw -q add 00100 allow ip6 from any to any
/sbin/ipfw -q add 00101 allow icmp6 from any to any

## Include file for customizations
## USE WITH CAUTION
incfile=/usr/local/nessy2/etc/ipfw/ipfw_open.rules.inc
if [ -f $incfile ]; then
    . $incfile
fi
```



ipfw_open.rules.inc

```
/sbin/ipfw -q add 5 deny icmp from any to any
echo "PING requests are blocked"
```

Figure M.2. Example of non-supported firewall rules in open mode

Before configuring non-supported rules, keep in mind that there is no validation check for these files, so **if you misconfigure either you can lose access to your appliance**, via SSH or otherwise.

To incorporate non-supported firewall rules in open mode

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/ipfw/`, edit the file `ipfw_open.rules` to incorporate the non-supported firewall rules of your choice.

Keep in mind that if you misconfigure this file you can lose access to your appliance.

4. Restart the firewall daemon to take into account your changes with the command:

```
service ipfw restart
```

Configuring Non-Supported Apache Settings

You can configure non-supported Apache settings that the configuration file `httpd` automatically incorporates:

- Before any existing configuration, thanks to the settings specified in the file(s) in the directory `/pre`.
- After all the existing configurations, thanks to the settings specified in the file(s) in the directory `/post`.

By default, both directories contain a file `httpd.conf.inc` that you can edit to specify the settings of your choice. You can also add as many `*.inc` files as you need in these directories.

httpd.conf

```
## Include files for customizations
## USE WITH CAUTION
IncludeOptional /usr/local/nessy2/etc/httpd/pre/*.inc

#
# This is the main Apache HTTP server configuration file. It contains the
# configuration directives that give the server its instructions.
...
ServerRoot "/usr/local"
...
Listen 80

LoadModule authn_file_module libexec/apache24/mod_authn_file.so
LoadModule authn_core_module libexec/apache24/mod_authn_core.so
LoadModule authz_host_module libexec/apache24/mod_authz_host.so
...

# Secure (SSL/TLS) connections
#include etc/apache24/extra/httpd-ssl.conf
#
# Note: The following must must be present to support
# starting without SSL on platforms with no /dev/random equivalent
# but a statically compiled-in mod_ssl.
#
<IfModule ssl_module>
SSLRandomSeed startup builtin
SSLRandomSeed connect builtin
</IfModule>

Include etc/apache24/Includes/*.conf

## Include files for customizations
## USE WITH CAUTION
IncludeOptional /usr/local/nessy2/etc/httpd/post/*.inc
```



myhttpconf.inc

```
Listen 81
```



httpd.conf.inc

```
<VirtualHost 10.10.10.10:81>
DocumentRoot /data1/customer-portal
ServerName portal.customer.com
ErrorLog "/var/log/httpd-portal_error_log"
<Directory /data1/customer-portal>
AllowOverride All
Order deny,allow
allow from all
</Directory>
</VirtualHost>
```

Figure M.3. Example of non-supported Apache settings

Before configuring non-supported settings, keep in mind that there a validation check for this file but **any invalid configuration may prevent the service from running properly or prevent you from accessing the GUI altogether.**

To incorporate non-supported Apache settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and *root* credentials.
3. Incorporate non-supported configurations in the file that suits your needs:

- Before any other instructions, edit the file *httpd.conf.inc* according to your needs. The full path to the file is */usr/local/nessy2/etc/httpd/pre/httpd.conf.inc*

You can also add any *.inc* file in the directory.

- After all instructions, edit the file *httpd.conf.inc* according to your needs. The full path to the file is */usr/local/nessy2/etc/httpd/post/httpd.conf.inc*

You can also add any *.inc* file in the directory.

4. Make sure the whole configuration file is still viable using the command:

```
apachectl configtest
```

If no errors are returned and the file syntax is OK, go to the next step. If not, you must edit the content of the included file(s) because you might no longer be able to access or GUI or prevent the service from running.

5. Once the configuration is OK, restart the Apache daemon to take into account your changes with the command:

```
apachectl restart
```

6. Make sure the daemon is running with the command:

```
apachectl status
```

Configuring Non-Supported Unbound Settings

You can configure non-supported Unbound settings that the configuration file `unbound` automatically incorporates:

- In the section `server` of the configuration file.
- In the section `remote-control` of the configuration file.

unbound.conf

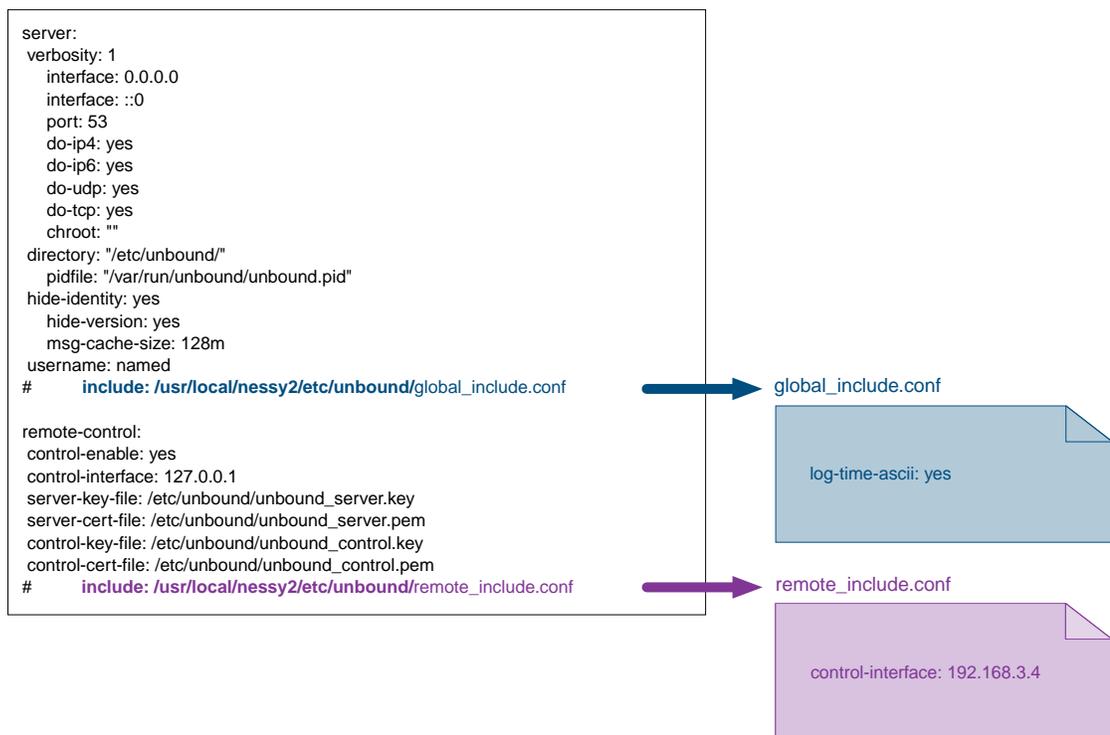


Figure M.4. Example of non-supported Unbound settings

Before configuring non-supported settings, keep in mind any invalid option is ignored.

To incorporate non-supported Unbound settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. Incorporate non-supported options in the file that suits your needs:
 - To configure non-supported server settings, edit the file `global_include.conf` according to your needs. The full path to the file is `/usr/local/nessy2/etc/unbound/global_include.conf`
 - To configure non-supported remote management settings, edit the file `remote_include.conf` according to your needs. The full path to the file is `/usr/local/nessy2/etc/unbound/remote_include.conf`
4. Make sure the whole configuration file is still viable using the command:

```
unbound-checkconf /etc/unbound/unbound.conf
```

If no errors are returned, go to the next step. If not, you must edit the content of the included file(s) because incorrect configurations are ignored.

5. Once the configuration is OK, restart the Unbound daemon to take into account your changes with the command:

```
service ipmdns.sh restart
```

6. Make sure the daemon is running with the command:

```
service ipmdns.sh status
```

Configuring Non-Supported NSD Settings

You can configure non-supported Unbound settings that the configuration file `nsd` automatically incorporates:

- In the section `remote-control` of the configuration file.
- In the section `server` of the configuration file.

`nsd.conf`

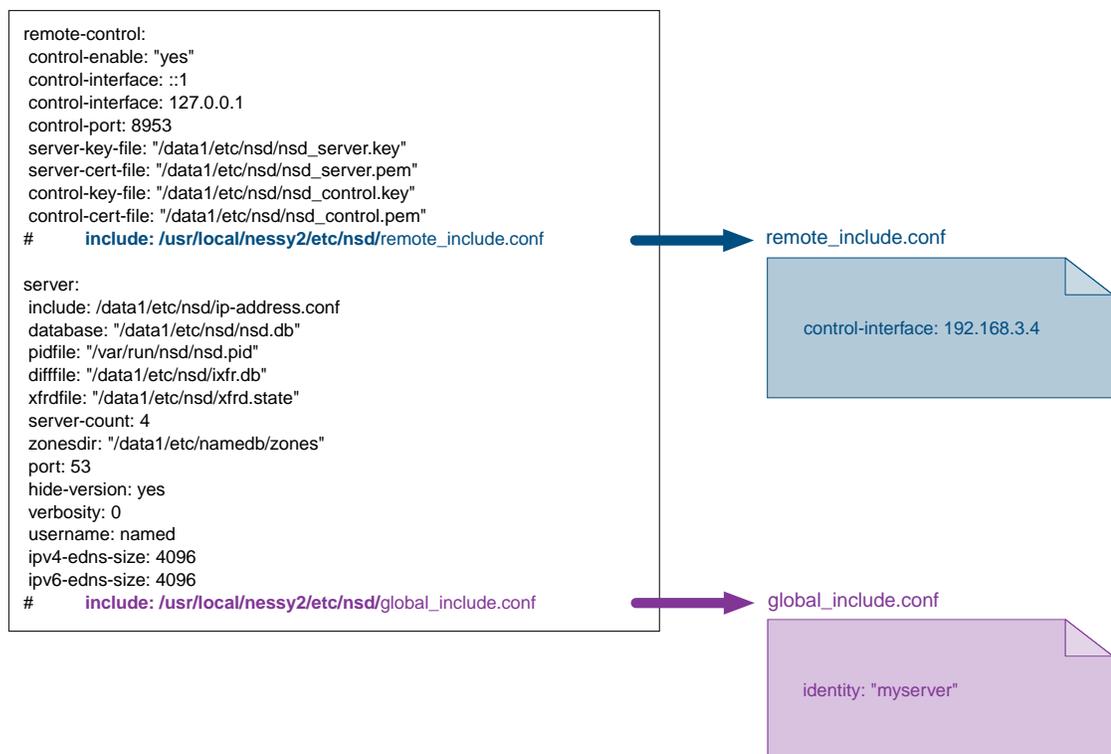


Figure M.5. Example of non-supported NSD settings

Before configuring non-supported settings, keep in mind any invalid option is ignored.

To incorporate non-supported NSD settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. Incorporate non-supported options in the file that suits your needs:
 - To configure non-supported remote management settings, edit the file `remote_include.conf` according to your needs. The full path to the file is `/usr/local/nessy2/etc/nsd/remote_include.conf`
 - To configure non-supported server settings, edit the file `global_include.conf` according to your needs. The full path to the file is `/usr/local/nessy2/etc/nsd/global_include.conf`
4. Make sure the whole configuration file is still viable using the command:

```
nsd-checkconf /etc/nsd/nsd.conf
```

If no errors are returned, go to the next step. If not, you must edit the content of the included file(s) because incorrect configurations are ignored.

5. Once the configuration is OK, restart the NSD daemon to take into account your changes with the command:

```
service ipmdns.sh restart
```

6. Make sure the daemon is running with the command:

```
service ipmdns.sh status
```

Configuring Non-Supported BIND Settings

You can configure non-supported BIND options that the configuration file `named` automatically incorporates:

- At server level, non-supported settings can be incorporated in the sections *global* and *options* of the configuration file.
- At view level, non-supported settings must be manually incorporated in the clause of each view declared in the server configuration file.

Before configuring non-supported settings, keep in mind any invalid option is ignored.

Configuring Non-Supported BIND Settings on a Server

At server level, you can incorporate non-supported settings and parameters to a BIND configuration file thanks to two files:

`global_include.conf`

Allows to incorporate statements like logging, masters, server, trusted-keys or managed-keys. In this file, **you must specify each statement** and its value.

`options_include.conf`

Allows to incorporate options. In this file, **no need to specify the statement "options"** because all the values that you add are automatically part of the statement.

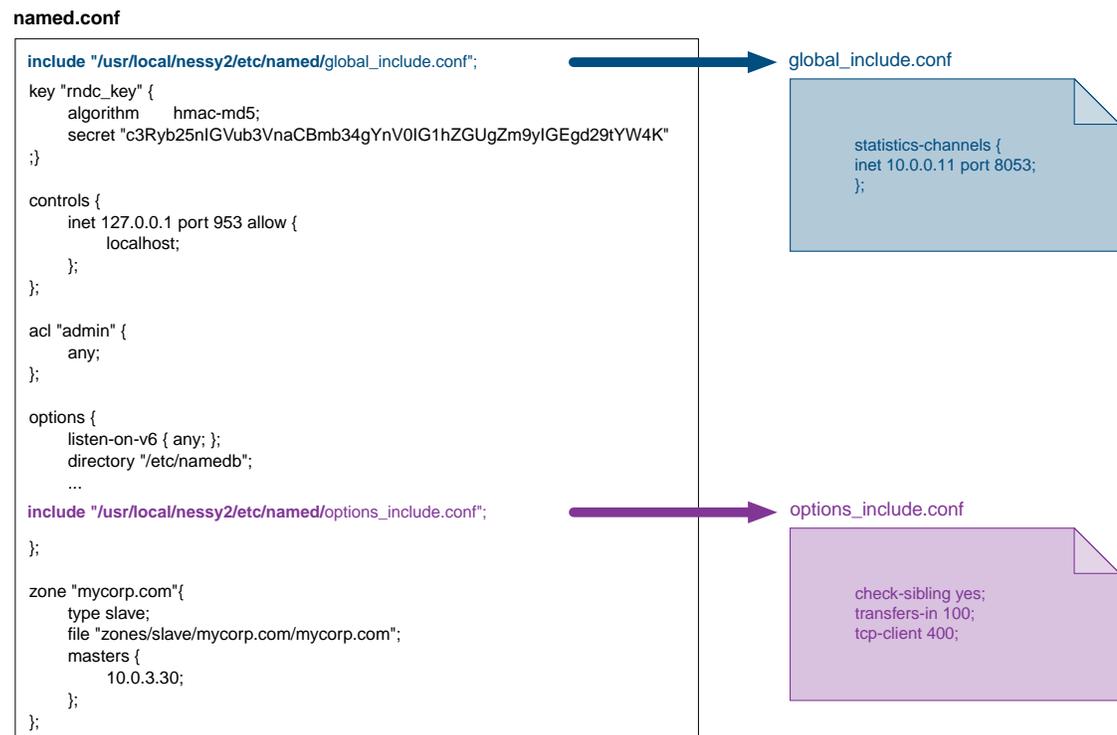


Figure M.6. Example of non-supported BIND settings configured for a server

Before configuring non-supported settings, keep in mind any invalid option is ignored.

To incorporate non-supported BIND settings on a server

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance or Linux server using an SSH session or a port console and *root* credentials.
3. Incorporate non-supported BIND settings for the server in the file that suits your needs:
 - To configure settings in the section *global*, edit in the file *global_include.conf* according to your needs. The full path to the file is */usr/local/nessy2/etc/named/global_include.conf*
 - To configure settings in the section *options*, edit in the file *options_include.conf* according to your needs. The full path to the file is */usr/local/nessy2/etc/named/options_include.conf*

4. Make sure the whole configuration file is still viable using the command:

```
/usr/local/nessy2/bin/named-checkconf /etc/namedb/named.conf
```

If no errors are returned and the configuration file is OK, go to the next step. If not, you must edit the content of the included file because incorrect configurations are ignored.

5. Once the configuration is OK, restart the DNS daemon to take into account your changes with the command:

```
service ipmdns.sh restart
```

If you installed Linux packages, you must stop and start the daemon using the commands:

```
service ipmdns stop service ipmdns start
```

6. Make sure the DNS daemon is running with the command:

```
service ipmdns.sh status
```

If you installed Linux packages, you must use the following command:

```
service ipmdns status
```

Configuring Non-Supported BIND Settings on a View

At view level, you can incorporate non-supported settings and parameters to a BIND configuration file if you create the file:

view_<dnsview-name>_include.conf

Allows to incorporate statements like *empty-zones-enable* or *cleaning-interval* to a clause "view". In this file, **no need to specify the statement "view"**, all the values that you add are automatically taken into account.

named.conf

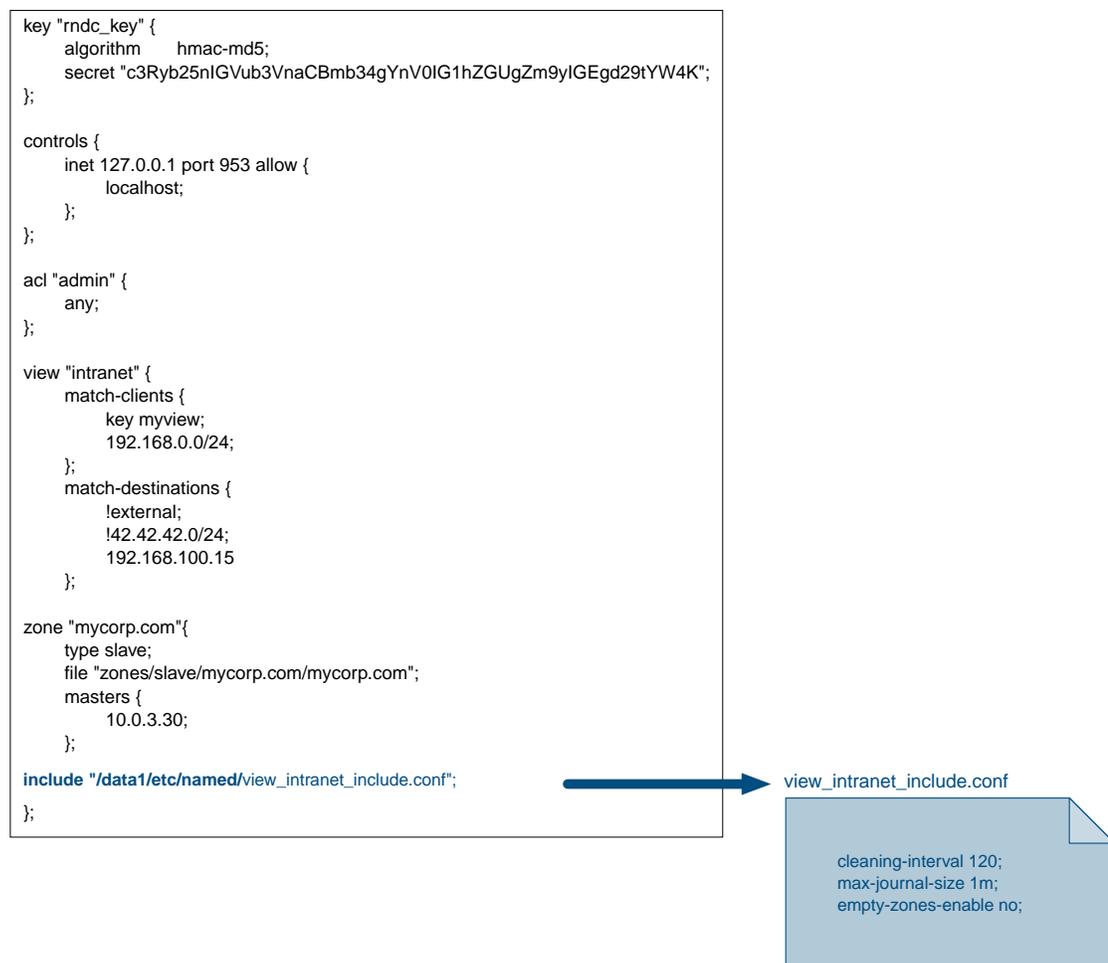


Figure M.7. Example of non-supported BIND settings configured for a view called "intranet"

Before configuring non-supported settings, keep in mind any invalid option is ignored.

To incorporate non-supported BIND settings on a view

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance or Linux server using an SSH session or a port console and *root* credentials.

3. In the directory `/data1/etc/namedb/`, create the file `view_<dnsview-name>_include.conf`, where `<dnsview-name>` is the name of one of your existing views, and incorporate the non-supported BIND settings of your choice for the view.

4. Make sure the whole configuration file is still viable using the command:

```
/usr/local/nessy2/bin/named-checkconf /etc/namedb/named.conf
```

If no errors are returned and the configuration file is OK, go to the next step. If not, you must edit the content of the included file because incorrect configurations are ignored.

5. Once the configuration is OK, restart the DNS daemon to take into account your changes with the command:

```
service ipmdns.sh restart
```

If you installed Linux packages, you must stop and start the daemon using the commands:

```
service ipmdns stop service ipmdns start
```

6. Make sure the DNS daemon is running with the command:

```
service ipmdns.sh status
```

If you installed Linux packages, you must use the following command:

```
service ipmdns status
```

Keep in mind that if any option is invalid, the included configuration is ignored until you correct what needs to be changed. As for conflicting options, they overwrite your configuration.

Configuring Non-Supported SNMP Settings

You can configure non-supported SNMP settings that the configuration file `snmpd` automatically incorporates after all the directives of the configuration file.

By default, the directory `/snmpd` contains a file `custom.conf` that can be edited to specify the settings of your choice. You can also add as many files as you need in this directory.

snmpd.conf

```
sysDescr EfficientIP SOLIDserver
sysObjectID .1.3.6.1.4.1.2440
sysServices 72
sysContact EfficientIP Support <support@efficientip.com>
sysLocation Unknown
master agentx
agentaddress udp:161
com2sec secu0 default 'public'
group MyROGroup v1 secu0
group MyROGroup v2c secu0
view all included .1 80
access MyROGroup "" any noauth exact all none none
access MyRWGroup "" any noauth exact all none none

## EIP specific conf
includeFile /data1/share/snmp/eip-snmpd.conf
includeFile /data1/share/snmp/eip-traps.conf

agentXPerms 0775 0775 root agentx
## Include file for customizations
## USE WITH CAUTION
includeDir /usr/local/nessy2/etc/snmpd
```



custom.conf

```
trapcommunity public
trapsess -v 2c -c public 10.0.11.3
authtrappable 1
```

Figure M.8. Example of non-supported SNMP settings

Before configuring non-supported settings, keep in mind that **there is no validation check for this file**. Any invalid configuration may prevent the service from running properly.

To incorporate non-supported SNMP settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/snmpd`, edit the file `custom.conf` according to your needs or create a file to incorporate the non-supported SNMP options of your choice.
4. Once the configuration is OK, restart the SNMP daemon to take into account your changes with the command:

```
service snmpd restart
```

5. Make sure the daemon is running with the command:

```
service snmpd status
```

Configuring Non-Supported DHCP Configurations

You can configure non-supported DHCP configurations that the configuration files `dhcpcd` and `dhcpcd6` automatically incorporates.

Before configuring non-supported options, keep in mind any invalid option is ignored.

Configuring Non-Supported DHCP Configurations in IPv4

By default, the directory `/dhcp` contains a file `global_include.conf` that can be edited to specify the configurations of your choice.

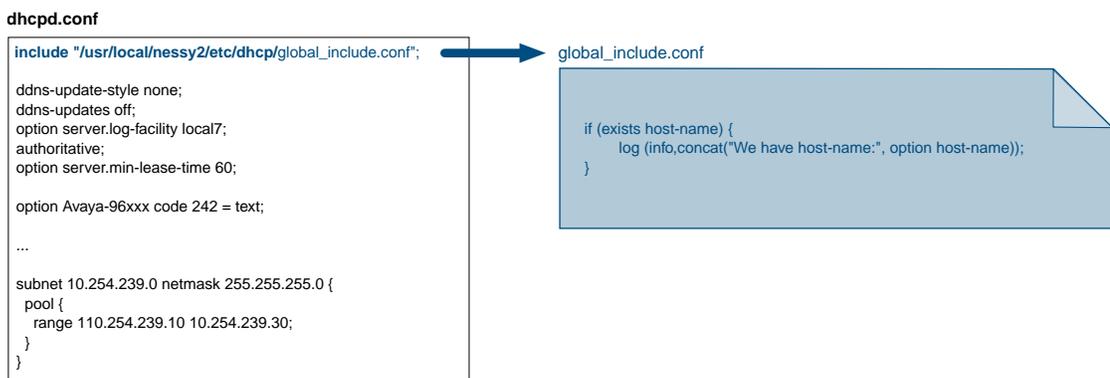


Figure M.9. Example of a non-supported DHCP configuration

Before configuring non-supported options, keep in mind any invalid option is ignored.

To incorporate non-supported DHCP configurations

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/dhcp`, edit the file `global_include.conf` according to your needs to incorporate the non-supported DHCP configurations of your choice.
4. Make sure the whole configuration file is still viable using the command:

```
/usr/local/nessy2/bin/dhcpd -t -cf /data1/etc/dhcpd.conf
```

If no errors are returned, go to the next step. If not, you must edit the content of the included file(s) because incorrect configurations are ignored.

5. Once the configuration is OK, restart the DHCP daemon to take into account your changes with the command:

```
service ipmdhcp.sh restart
```

6. Make sure the daemon is running with the command:

```
service ipmdhcp.sh status
```

Configuring Non-Supported DHCP Configurations in IPv6

By default, the directory `/dhcp6` contains a file `global_include.conf` that can be edited to specify the configurations of your choice.

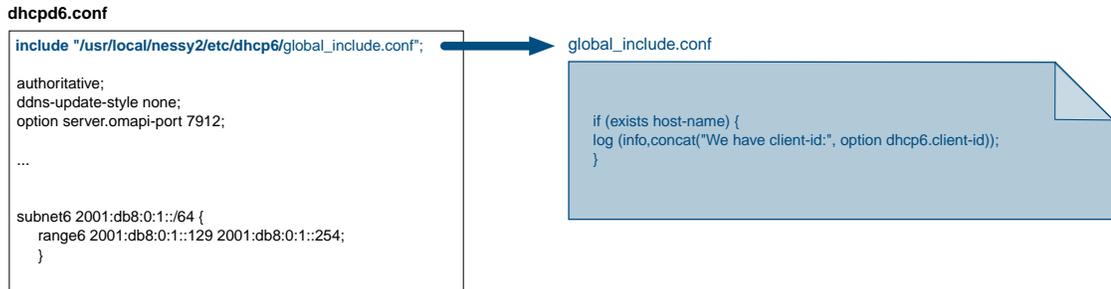


Figure M.10. Example of a non-supported DHCPv6 configuration

Before configuring non-supported options, keep in mind any invalid option is ignored.

To incorporate non-supported DHCPv6 configurations

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and *root* credentials.
3. In the directory `/usr/local/nessy2/etc/dhcp6`, edit the file `global_include.conf` according to your needs to incorporate the non-supported DHCP configurations of your choice.
4. Make sure the whole configuration file is still viable using the command:

```
/usr/local/nessy2/bin/dhcpd -6 -t -cf /data1/etc/dhcpd6.conf
```

If no errors are returned, go to the next step. If not, you must edit the content of the included file(s) because incorrect configurations are ignored.

5. Once the configuration is OK, restart the NSD daemon to take into account your changes with the command:

```
service ipmdhcp6.sh restart
```

6. Make sure the daemon is running with the command:

```
service ipmdhcp6.sh status
```

Configuring Non-Supported NTP Settings

You can configure non-supported NTP settings that the configuration file `ntp` automatically incorporates.

By default, the directory `/etc` contains a file `ntp.conf.inc` that can be edited to specify the settings of your choice.

ntp.conf

```
# By default, exchange time with everybody, but dont allow configuration.
restrict -4 default kod notrap nomodify nopeer noquery limited
restrict -6 default kod notrap nomodify nopeer noquery limited

# Local users may interrogate the ntp server more closely.
restrict 127.0.0.1
restrict ::1

## Include file for customizations
## USE WITH CAUTION
includefile /usr/local/nessy2/etc/ntp.conf.inc
```



ntp.conf.inc

```
peer 192.168.5.6 key 11
peer 2001:db8:1:100 key 33
keys /data1/exports/ntp.keys
trustedkey 11 33
```

Figure M.11. Example of non-supported NTP settings

Before configuring non-supported settings, keep in mind that **there is no validation check for this file**. Any invalid configuration may prevent the service from running properly.

To incorporate non-supported NTP settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/`, edit the file `ntp.conf.inc` according to your needs to incorporate the non-supported NTP options of your choice.
4. Once the configuration is OK, restart the NTP daemon to take into account your changes with the command:

```
service ntpd restart
```

5. Make sure the daemon is running with the command:

```
service ntpd status
```

Configuring Non-Supported syslog-ng Settings

You can configure non-supported NTP settings that the configuration file `syslog-ng` automatically incorporates.

By default, the directory `/include` can contain any file that you can configure with the settings of your choice.

syslog.conf

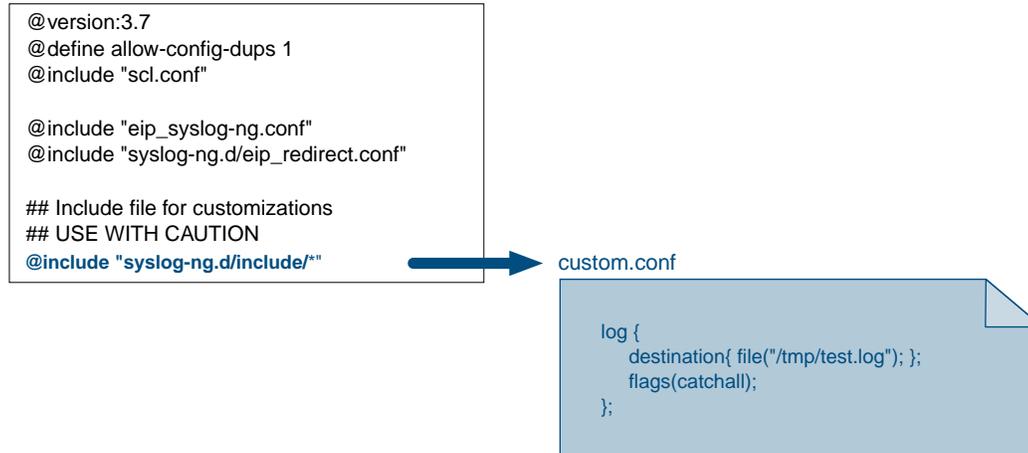


Figure M.12. Example of non-supported syslog-ng settings

Before configuring non-supported options, keep in mind that:

- **There is no validation check for these files.** Any invalid configuration may prevent the service from running properly.
- The user definitions can override the default configuration.

To incorporate non-supported syslog-ng options

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/syslog-ng.d/include/`, edit the file `custom.conf` according to your needs to incorporate the non-supported syslog-ng settings of your choice.
4. Reload the configuration file to take into account the changes using the command:

```
syslog-ng-ctl reload
```

If no errors are returned, go to the next step. If not, you must edit the content of the included files because incorrect configurations are ignored.

5. Make sure the daemon is running with the command:

```
service syslog-ng status
```

Configuring Non-Supported PostgreSQL Settings

You can configure non-supported NTP settings that the configuration file `postgresql` automatically incorporates.

postgresql.conf

```
log_destination = 'syslog'
autovacuum = on
max_connections = 100
maintenance_work_mem = 64MB
update_process_title = off
wal_level = hot_standby
max_wal_senders = 2
listen_addresses = '*'
hot_standby = on
wal_keep_segments = 256
log_min_duration_statement = -1
synchronous_commit = off
commit_delay = 50000
commit_siblings = 5
wal_writer_delay = 2000
geqo_threshold = 9

## Include file for customizations
## USE WITH CAUTION
include_if_exists 'usr/local/nessy2/etc/postgresql.conf.inc'
```

postgresql.conf.inc

```
shared_preload_libraries = 'auto_explain'
auto_explain.log_min_duration = '1000ms'
```

Figure M.13. Example of non-supported PostgreSQL settings

Before configuring non-supported options, keep in mind that **there is no validation check for these files**. Any invalid configuration may prevent the service from running properly.

To incorporate non-supported PostgreSQL settings

1. Meet the [Prerequisites](#) and take into account the [Limitations](#).
2. Connect to your appliance using an SSH session or a port console and `root` credentials.
3. In the directory `/usr/local/nessy2/etc/`, edit the file `postgresql.conf.inc` according to your needs to incorporate the non-supported PostgreSQL options of your choice.
4. Restart the PostgreSQL daemon to take into account your changes with the command:

```
service postgresql restart
```

5. Make sure the daemon is running with the command:

```
service postgresql status
```

Index

A

ACL

- creating an ACL based on option 82, 416
- DNS server ACL, 519
- editing a view match clients list, 542
- editing a view match destination list, 542
- granting access to known clients, 397
- range ACL, 383
- restricting access, 397

Active Directory

- authenticating administrators, 1040
- domain, 578
- dynamic update authentication, 635
- relying on AD credentials to log users, 1038

addresses, 247

- adding, 249
- adding by search, 252
- adding manually, 249
- assigning, 249
- configuring aliases, 254
- deleting, 262
- editing, 253
- editing aliases, 256
- editing the network/broadcast address, 253
- moving IPv4 addresses across networks, 260, 261
- moving IPv4 addresses across spaces, 260
- moving IPv4 addresses across the VLSM, 293
- pinging an address, 262
- removing aliases, 256
- renaming Pv4 addresses massively, 259
- restoring, 263
- statuses and types, 248
- undoing an address deletion, 263
- updating

- Device Manager, 963

Administration

- importing CSV data, 159

administrators

- authenticating
 - Active Directory, 1040
 - LDAP, 1042
 - RADIUS, 1043
- authenticating remote users, 1038
- local user, 1030
- managing group of users, 1019
- managing users, 1030

advanced properties, 732

- configuring
 - DHCP, 746
 - DNS, 748

- IPAM, 735

- inheriting, 750
- internal module setup, 26
- propagating, 750
- restricting, 750
- setting, 750

Agentless server

- DHCP (see Microsoft DHCP)
- DNS (see Microsoft DNS)

alerts, 1088

- adding alerts, 1090
- checking alerts, 1093
- disabling alerts, 1093
- dismissing alerts, 1094
- editing alerts, 1092
- enabling alerts, 1093

allow-notify

- at server level, 499
- at view level, 548
- at zone level, 599

allow-query

- at server level, 501
- at view level, 551
- at zone level, 602

allow-query-cache

- at server level, 502
- at view level, 552

allow-recursion, 497

allow-transfer

- at server level, 503
- at view level, 554
- at zone level, 603

allow-update, 604

also-notify

- at server level, 499
- at view level, 548

alt-transfer-source, 515

alt-transfer-source-v6, 515

Amazon Route 53

- DNS servers, 482

analytics

- DHCP data sampling, 424
- DNS data sampling, 713
- Guardian, 810

answerlog, 721

anycast

- for BGP, 527
- for IS-IS, 532
- for OSPF, 523

appliance

- default gateway, 84
- high availability (see high availability management)
- reboot, 1129
- remote management, 1048

- shutdown, 1130
- SNMP agent, 107
- time and date configuration, 79
- troubleshooting, 1120
- upgrading, 1133
- Application, 752
 - applications, 756
 - adding and deploying, 757
 - associating with GSLB servers, 761
 - deleting, 762
 - dissociating from GSLB servers, 761
 - editing, 761
 - configuring the module, 753
 - labels in IPv6, 1161
 - limitations, 753
 - nodes, 768
 - adding, 770
 - deleting, 774
 - editing, 773
 - managing/unmanaging, 774
 - pools, 763
 - adding, 764
 - deleting, 766
 - editing, 766
 - prerequisites, 753
 - rights, 1023
 - service configuration, 105
 - traffic policy, 756
 - adding and deploying, 759
- autoselect media port, 97
- AWS
 - Amazon Route 53 DNS servers, 482
- B**
- backup, 1123
 - configuring the remote archive, 1126
 - creating an instant backup, 1124
 - restoring a backup, 1127
 - scheduling a daily backup, 1125
 - setting a backup rotation, 1125
 - uploading a backup, 1127
- BGP (see anycast)
- BIND
 - DNS servers, 468
 - importing zones archive file, 175
- blackhole, 504
- bookmarks, 74
- C**
- CARP (Common Address Redundancy Protocol), 93
- centralized management, 1048
- certificate
 - choosing a certificate, 104
 - creating a certificate, 1107
 - creating a CSR, 1109
 - creating a private key, 1111
 - HTTPS, 104
 - importing a certificate, 1106
 - importing a CSR file, 1106
 - importing a private key, 1106
- Class Studio, 1164
 - class objects
 - adding, 1176
 - deleting, 1221
 - editing, 1220
 - organizing, 1221
 - classes
 - adding, 1168
 - applying, 1169
 - deleting, 1173
 - disabling from a listing page, 1173
 - disabling from Class Studio, 1173
 - duplicating, 1171
 - enabling from a listing page, 1173
 - enabling from Class Studio, 1173
 - moving, 1172
 - renaming, 1170, 1171
 - using another or no class, 1172
 - rights, 1023
 - syntax, 1221
- classes (see Class Studio)
- configuration file
 - DHCP (dhcpd.conf), 108
 - DHCPv6 (dhcpd6.conf), 108
 - DNS (named.conf), 108
 - NSD (nsd.conf), 108
 - Unbound (unbound.conf), 108
- Custom DB, 1223
 - custom data
 - adding, 1225
 - deleting, 1226
 - editing, 1226
 - importing CSV data, 162
 - custom database
 - adding, 1224
 - deleting, 1225
 - editing, 1224
 - importing custom data through a CSV file, 162
- customize
 - add image to login page, 1151
 - add image to welcome banner, 1153
 - edit welcome banner, 1153
 - hide welcome banner, 1154
 - meta database, 1223
 - set your own field names, 1155

D

- Dashboards, 189
 - dashboards, 190
 - assigning a gadget, 191
 - hiding gadgets, 192
 - organizing gadgets, 191
 - gadgets, 193
 - assigning, 211
 - creating, 207
 - default gadgets, 1243
 - deleting, 219
 - displayed by default, 198
 - editing, 214
 - types, 195
 - visibility, 216
- DDNS, 635
 - (see also secure dynamic update)
 - server configuration, 517
 - zone authorizations, 604
- Device Manager, 937
 - devices, 938
 - adding automatically, 940
 - adding manually, 944
 - deleting, 947
 - duplicating, 946
 - managing status and visibility, 939
 - merging, 946
 - dual stack, 963
 - importing CSV data, 147
 - interfaces, 948
 - adding automatically, 950
 - adding manually, 954
 - deleting, 961
 - editing, 959
 - managing status and visibility, 949
 - renaming, 956
 - IPAM interaction, 963
 - using advanced properties, 966, 969
 - using the link option, 967
 - using the MAC address, 963
 - monitoring changes, 960
 - ports, 948
 - adding automatically, 950
 - adding manually, 952
 - deleting, 961
 - editing, 957
 - managing status and visibility, 949
 - renaming, 956
 - rules, 971
 - from the DHCP, 971
 - updating from
 - IPAM addresses automatically, 943
 - IPAM addresses manually, 963
 - NetChange (devices), 942
 - NetChange (ports & interfaces), 950
 - devices (see Device Manager)
- DHCP, 303
 - ACL (see ACL)
 - groups, 376
 - importing
 - Infoblox configuration file, 172
 - ISC configuration file, 167
 - MetalP configuration file, 173
 - Microsoft configuration file, 171
 - NetID, 173
 - VitalQIP configuration file, 169
 - importing CSV data, 134
 - IP helper, 364
 - labels in IPv6, 1161
 - leases (see leases)
 - monitoring
 - audit, 424
 - from the page Analytics, 424
 - from the properties page, 423
 - lease statistics, 423
 - server analytics, 424
 - state log, 424
 - using rules, 430
 - ping check, 401
 - preventing IP address duplication, 401
 - ranges (see ranges (DHCP))
 - relay agents, 364
 - relay agent information, 395
 - relay agent information (DHCPv6), 396
 - reporting (see reports)
 - rights, 1023
 - scopes (see scopes)
 - servers, 332
 - EfficientIP servers, 334
 - ISC DHCP server, 342
 - Microsoft servers, 337
 - statuses, 333
 - smart architectures, 305
 - SSL vs SNMP, 334
 - statics (see statics)
- DHCP options, 410
 - advanced configuration, 1245
 - bulk changes at range level, 386
 - bulk changes at scope level, 358
 - bulk changes at static level, 374
 - circuit ID, 415
 - DHCPv4, 412
 - DHCPv6, 413
 - group, 378
 - range, 386
 - relay agent information, 415
 - relay agent information (DHCPv6), 417

- remote ID, 415
- RFC 2132/option 43, 417
- RFC3046/option 82, 415
- RFC3315/DHCPv6, 415
- scope, 358
- server, 352
- setting options, 411
- static, 373
- vendor class identifier, 415
- vendor specific information, 417
- DHCPv6 prefix delegation, 420
 - adding prefixes, 421
 - browsing prefixes, 420
 - deleting prefixes, 422
- DNS, 433
 - anycast, 523
 - blackhole, 504
 - configuring a resolver, 86
 - configuring servers, 494
 - DNS keys, 521
 - DNS64, 510
 - configuring, 512
 - disabling, 514
 - supported substatements, 511
 - DNSCrypt, 535
 - DNSSEC (see DNSSEC)
 - EDNS options
 - at server level, 506
 - at view level, 556
 - forwarding
 - server, 494
 - view, 546
 - zone, 598
 - Guardian (see Guardian)
 - Hybrid (see Hybrid)
 - importing
 - BIND archive file, 175
 - VitalQIP archive file, 177
 - importing CSV data, 143
 - including non-supported settings, 523
 - minimal-responses, 507
 - monitoring
 - answerlog, 721
 - audit, 713
 - from the page Analytics, 713
 - from the properties page, 712
 - queries and answers, 719
 - query statistics, 712
 - querylog, 719
 - state log, 713
 - recursion at server level, 496
 - recursion at view level, 549
 - reporting (see reports)
 - resolver, 461
 - resource records (see resource records)
 - rights, 1022
 - RPZ (see RPZ)
 - RRL, 508
 - servers, 459
 - Amazon Route 53 servers, 482
 - BIND servers, 468
 - EfficientIP servers, 461
 - generic servers, 478
 - Microsoft servers, 465
 - Nominum servers, 480
 - statuses, 461
 - synchronizing, 490
 - smart architectures, 435
 - sortlist, 507
 - at view level, 556
 - sources, 514
 - at view level, 557
 - at zone level, 606
 - views (see views)
 - zone delegation, 596, 630
 - zones (see zones)
- DNS firewall (see RPZ)
- DNS64, 510
- DNSSEC, 676
 - chain of trust, 694
 - disabling
 - on authoritative servers, 693
 - on recursive servers, 699
 - signing keys, 691
 - enabling
 - signing keys, 692
 - generating a newKSK, 687
 - keys
 - KSK, 684
 - trust anchor, 696
 - ZSK, 684
 - on authoritative servers, 677
 - on recursive servers, 694
 - records
 - DNSKEY, 677
 - DS, 678
 - NSEC/NSEC3, 678
 - NSEC3PARAM, 678
 - publishing the DS, 682
 - RRSIG, 678
 - resolvers, 694
 - rollover
 - KSK, 686
 - ZSK, 685
 - signing a zone, 679
 - unsigned zones, 693
- domains (VLAN), 974
 - adding, 975

- deleting, 976
- editing, 975
- Dual stack
 - Device Manager, 963
 - interfaces, 83
 - virtual interface, 96
- dynamic update, 635
 - (see also secure dynamic update)
 - at zone level, 604
- E**
- EDNS
 - on a DNS server, 506
 - on a DNS view, 556
- EfficientIP
 - DHCP servers, 334
 - DNS servers, 461
- exports, 178
 - configuring exports, 180
 - export files, 186
 - export templates, 187
 - finding exports in the database, 180
 - reimport, 182
 - scheduled export configuration files, 187
- F**
- failover, 403
 - communications-interrupted state, 404
 - configuring virtual IP, 94
 - DHCP Safe Failover, 403
 - Ethernet port, 92
 - management database, 1058
 - normal state, 404
 - partner-down state, 405
 - switching to partner-down, 408
- Farm, 448
- file transfer
 - using SFTP/SCP/RSYNC, 102
- firewall, 87
- forwarding
 - DNS (see DNS, forwarding)
 - logs, 1096
- G**
- generic
 - DNS servers, 478
- global search, 35
- group of users, 1019
 - adding groups, 1020
 - assigning resources, 1023
 - defining a group rights, 1026
 - deleting groups, 1029
 - editing groups, 1028
 - enabling and disabling groups, 1028
 - importing CSV data, 160
- GSS-TSIG, 635
 - configuring a server, 639
 - configuring a zone, 640
 - generating a GSS-TSIG key, 638
 - uploading a GSS-TSIG key, 638
- Guardian, 775, 776
 - cache, 784
 - clearing automatically, 789
 - clearing manually, 788
 - forcing entries expiration, 788
 - resetting, 786
 - restoring, 787
 - saving, 787
 - sending, 792
 - sharing, 790
 - configuration, 779
 - configuring from the GUI, 105
 - enabling, 777
 - limitations, 776
 - monitoring from the GUI, 807
 - analytics, 810
 - statistics, 807
 - prerequisites, 776
 - protection, 820
 - client log, 846
 - disabling, 849
 - enabling, 820
 - lists, 824
 - policies, 833
 - rescue mode, 821
 - triggers, 836
 - views, 829
 - rights, 1023
 - statistics, 793
 - client statistics, 802
 - server statistics, 793
- GUI, 28
 - bookmarks, 74
 - breadcrumb, 38
 - charts, 61
 - contextual menu, 58
 - customizing welcome banner and login page, 1151
 - global search, 35
 - home page, 29
 - listing pages, 44
 - customizing the list layout, 52
 - filtering data, 47
 - listing templates, 52
 - sorting data, 46
 - Main Dashboard, 29
 - menu, 40
 - modules, 30

- multi-status (column), 58
- notifications, 37
- properties pages, 60
- quick wizards, 71
- sidebar, 29
- top bar, 35
- tree view, 32
- wizards, 65

H

- high availability management, 1058
 - adding a remote appliance, 1052
 - configuring, 1060
 - configuring the management appliance, 1051
 - controlling the automatic switch, 1064
 - disabling the configuration, 1075
 - editing the re-enrollment settings, 1066
 - enrolling the Hot Standby, 1060
 - frequently asked questions, 1073
 - hot standby appliance, 1049
 - keeping the database when disabling, 1071
 - limitations, 1059
 - master appliance, 1049
 - monitoring each appliance services, 1095
 - monitoring statistics, 1055
 - monitoring the logs, 1054
 - monitoring the time drift, 1055
 - network disruption, 1065
 - prerequisites, 1059
 - replacing an appliance, 1074
 - split-brain, 1070
 - standalone appliance, 1049
 - switching the appliances role, 1062
 - troubleshooting, 1068
 - updating the Hot Standby database, 1062
- HSM, 700
 - browsing the HSM servers, 701
 - disabling the HSM, 709
 - limitations, 701
 - prerequisites, 700
 - setting up the HSM, 703
- HTTPS
 - certificate, 104, 1106
- Hybrid, 667
 - backup, restoration and upgrade with, 675
 - checking the compatibility with, 667
 - forcing the compatibility with, 673
 - generating the incompatibilities report, 668
 - limitations, 673
 - NSD, 671
 - switching back to BIND, 674
 - switching to, 670
 - Unbound, 672

I

- imports
 - data
 - Administration, 159
 - CSV, 116
 - Device Manager, 147
 - DHCP, 134, 167
 - DNS, 143, 175
 - IPAM, 118, 164
 - NetChange, 146, 855
 - SPX, 156
 - VLAN Manager, 150
 - VRF, 153
 - from
 - BIND DNS, 175
 - Infoblox DHCP, 171
 - ISC DHCP, 167
 - MetalP DHCP, 172
 - Microsoft DHCP, 170
 - NetID DHCP, 173
 - NetID IPAM data, 165
 - VitalQIP DHCP, 168
 - VitalQIP DNS, 176
 - VitalQIP IPAM data, 164
 - import templates, 163
- Infoblox
 - importing a DHCP configuration file, 171
- interfaces (see Device Manager, interfaces)
 - Device Manager, 948
- IP addresses (see addresses)
- IPAM, 220
 - addresses (see addresses)
 - APNIC management (see SPX)
 - block-type networks (see networks, block-type)
 - importing
 - NetID host addresses, 166
 - NetID networks, 165
 - NetID subnets, 165
 - Vital QIP data, 164
 - importing CSV data, 118
 - IP addresses (see addresses)
 - labels in IPv6, 1161
 - networking, 220
 - pools (see pools)
 - provisioning, 271
 - rights, 1022
 - RIPE management (see SPX)
 - spaces (see spaces)
 - SPX management (see SPX)
 - subnet-type networks (see networks, subnet-type)
 - transition options, 265
 - updating
 - Device Manager (see addresses)

- Device Manager with addresses, 963
 - VLAN Manager with subnet-type networks, 984
- updating from
 - NetChange discovered items, 902
 - VLSM (see VLSM)
- IPLocator (see NetChange)
- IS-IS (see anycast)
- ISC
 - importing a DHCP configuration file, 167
 - ISC DHCP servers, 342
- K**
- keys
 - DNSSECsiging keys, 684
 - GSS-TSIG, 636
 - SSH (for SFTP backups), 1126
 - trust anchor, 696
 - TSIG
 - for dynamic updates at server level, 517
 - for dynamic updates at zone level, 604
 - to secure a server, 491
- L**
- labels
 - for IPv6 addressing, 1161
- lame-ttl
 - at server level, 505
 - at view level, 555
- LDAP
 - authenticating administrators, 1042
 - authentication for SSH, 1294
- leases, 388
 - blacklisting, 392
 - converting to statics, 391
 - lease time configuration, 389
 - pinging, 401
 - releasing lease, 391
 - tracking logs, 392
- license
 - activating a license, 112
 - adding a license, 26
 - exporting remote license keys, 111
 - installing a license, 112
 - renewing a license, 110
 - requesting a license, 25
 - requesting an license key, 110
- Linux
 - DHCP packages, 342
 - DNS packages, 468
- local files listing, 1112
- Locked Synchronization
 - on a DHCP smart architecture, 329
 - on a DNS smart architecture, 456
- logs
 - redirection, 1096
 - visualization, 1095
- M**
- maintenance, 1106
 - backup (see backup)
 - clearing SOLIDserver cache, 1120
 - local files listing (see local files listing)
 - maintenance mode, 1119
 - reboot (see appliance, reboot)
 - shutdown (see appliance, shutdown)
 - troubleshooting (see troubleshooting)
 - update macros and rules, 1120
- management
 - high availability (see high availability management)
- Master/Slave, 440
- max-cache-size
 - at server level, 505
 - at view level, 555
- MetalP
 - importing a DHCP configuration file, 172
- Microsoft DHCP
 - adding a server, 339
 - agentless server, 337
 - importing a configuration file, 170
 - limitations, 338
 - managing failover, 341
 - prerequisites, 338
- Microsoft DNS
 - adding a server, 467
 - agentless server, 465
 - limitations, 466
 - prerequisites, 465
- minimal-responses, 507
- monitoring, 1078
 - alerts (see alerts)
 - database tables size, 1104
 - DHCP data, 423
 - DNS data, 712
 - high availability appliances, 1054
 - netstat, 1104
 - remote appliances, 1054
 - reports (see reports)
 - services logs, 1095
 - services statistics, 1097
 - session tracking, 1099
 - SNMP (see SNMP, monitoring)
 - user tracking, 1099
 - copy user operations in syslog, 1100
- Multi-Master, 442

N

- NetChange, 850
 - Addresses, 897
 - configuration files versioning, 884
 - discovered items, 899
 - history, 903
 - purging, 902
 - refreshing, 901
 - updating the IPAM, 902
 - discovering network devices, 856
 - importing CSV data, 146
 - licenses, 851
 - monitoring
 - automating devices synchronization, 907
 - keeping the database up to date, 906
 - network devices, 852
 - adding, 854
 - connecting via a web console, 861
 - connecting via SSH, 861
 - connecting via telnet, 861
 - deleting, 862
 - making a snapshot, 861
 - refreshing, 859
 - scheduling a refresh, 860
 - statistics, 904
 - ports, 870
 - 802.1X authentication, 874
 - associating a port with a VLAN, 882
 - configuring tagging mode, 880
 - edge and terminal ports, 870
 - editing speed and duplex, 873
 - enabling/disabling, 872
 - interconnection ports, 870
 - limiting access using MAC addresses, 875
 - limiting access using user rights, 878
 - port-security, 875
 - refreshing, 883
 - statistics, 904
 - updating description, 874
 - rights, 1023
 - routes, 863
 - statistics, 904
 - tuning
 - customizing the devices' type, 908
 - updating
 - Device Manager with devices, 942
 - Device Manager with discovered items, 943
 - IPAM with addresses, 902
 - versioning, 884
 - comparing configuration files, 891
 - connection profiles, 886
 - disabling, 895
 - downloading configuration files, 893
 - enabling and configuring, 888
 - refreshing configuration files, 890, 894
 - VLANs, 867
 - adding, 868
 - deleting, 869
 - editing, 868
- Netchange
 - labels in IPv6, 1161
- NetID
 - importing a DHCP configuration file, 173
 - importing IPAM data, 165
- netstat, 1104
- network
 - adding a specific route, 85
 - basic interface configuration, 83
 - default gateway, 84
 - DNS resolver, 86
 - duplex, 97
 - Ethernet port failover, 92
 - interface IPv4/IPv6, 83
 - interface trunking (802.1q), 90
 - speed, 97
 - static route, 85
 - VHID, 94
 - virtual interface, 96
 - virtual IP, 94
- network flows
 - DHCP, 1236
 - DNS, 1237
 - High Availability, 1240
 - IPAM, 1235
 - NetChange, 1239
 - remote management, 1240
 - SOLIDserver, 1234
- networks
 - block-type
 - adding, 227
 - deleting, 241
 - editing, 234
 - moving, 237
 - network map, 238
 - splitting, 235
 - subnet-type
 - adding, 227
 - By search addition, 231
 - deleting, 241
 - discovering assigned addresses (IPv4 only), 237
 - editing, 234
 - editing the network/broadcast address, 253
 - finding available terminal networks, 231
 - managing/unmanaging, 239
 - merging, 236
 - moving, 237

- network map, 238
 - splitting, 235
 - statuses, 227
 - updating VLAN Manager, 984
 - Nominum
 - DNS servers, 480
 - notify
 - at server level, 499
 - at view level, 548
 - at zone level, 599
 - notify-source, 515
 - notify-source-v6, 515
 - NTP, 79
 - configuring the NTP server, 79
 - forcing an NTP update, 80
- O**
- One-to-Many, 314
 - One-to-One, 312
 - OSPF (see anycast)
- P**
- Packager, 1227
 - creating packages, 1228
 - deleting packages, 1231
 - downloading packages, 1231
 - installing packages, 1230
 - uninstalling packages, 1231
 - uploading packages, 1228
 - ping
 - addresses, 262
 - leases, 401
 - terminal networks (IPv4 only), 237
 - pools
 - adding, 244
 - deleting, 246
 - reserving, 245
 - resizing (IPv4 only), 245
 - port-security (see NetChange, ports)
 - ports (see Device Manager, ports)
 - Device Manager, 948
 - hardware appliance, 2
 - NetChange, 870
 - PXE, 400
 - changing the lease time for PXE clients, 401
 - next-server and filename options (v4), 400
- Q**
- query-source, 514
 - query-source-v6, 514
 - querylog, 719
 - quick wizards, 71
- R**
- RADIUS
 - authenticating administrators, 1043
 - configuring a FreeRadius server, 1289
 - configuring with Cisco RADIUS ACS, 1290
 - ranges (DHCP), 379
 - adding, 380
 - deleting, 387
 - editing, 383
 - options, 386
 - replicating in the IPAM, 385
 - resizing, 384
 - ranges (VLAN), 977
 - adding, 978
 - deleting, 980
 - editing, 978
 - resizing, 979
 - records
 - SOA
 - resetting, 578
 - recursion, 496
 - at server level, 496
 - at view level, 549
 - recursive-clients, 498
 - remote management, 1048
 - adding appliances, 1052
 - configuring the management appliance, 1051
 - deleting an appliances, 1075
 - monitoring statistics, 1055
 - monitoring the logs, 1054
 - monitoring the statistics, 1055
 - monitoring the time drift, 1055
 - prerequisites, 1052
 - remote network configuration, 1053
 - remote service configuration, 1053
 - replacing an appliance, 1074
 - upgrading remote appliances, 1135
 - reports, 1078
 - browsing the reports database, 1078
 - downloading and displaying, 1086
 - generating, 1084
 - managing scheduled reports, 1087
 - on DHCP scopes, 1079
 - on DHCP servers, 1079
 - on DNS servers, 1080
 - on DNS views, 1082
 - on DNS zones, 1082
 - on NetChange network devices, 1083
 - on statistics, 1084
 - on users, 1037
 - scheduling, 1085
 - resource records, 610
 - adding, 612

- A, 613
 - AAAA, 614
 - AFSDB, 614
 - CAA, 615
 - CERT, 615
 - CNAME, 616
 - DHCID, 617
 - DNAME, 617
 - DNSKEY, 618
 - DS, 684
 - HINFO, 619
 - MINFO, 619
 - MX, 620
 - NAPTR, 621
 - NS, 621
 - NSAP, 622
 - OPENPGPKEY, 623
 - PTR, 623
 - SRV, 625
 - SSHFP, 624
 - TLSA, 625
 - TXT, 626
 - URI, 626
 - WKS, 627
 - delegation, 630
 - deleting, 633
 - DNSSEC (see DNSSEC, records)
 - duplicating, 631
 - editing, 628
 - load balancing, 633
 - moving, 632
 - SOA, 610
 - SPF, 634
 - supported RRs, 610
 - Response policy zone (see RPZ)
 - Response Rate Limiting (RRL) (see DNS, RRL)
 - rights management, 1018
 - authentication rules, 1038
 - groups of users, 1019
 - users, 1030
 - route
 - adding a specific route, 85
 - static route, 85
 - RPZ, 646
 - limitations, 666
 - overriding rules, 664
 - policies
 - NODATA, 653
 - NXDOMAIN, 653
 - PASSTHRU, 653
 - REDIRECT, 652
 - poliicies (see RPZ, rules)
 - prerequisites, 646
 - records (see RPZ, rules)
 - rules, 652
 - based on domain names, 654
 - based on IP addresses, 657
 - based on name server domain name, 659
 - based on name server IP address, 659
 - based on specific resources, 663
 - deleting, 664
 - order, 654
 - zones
 - adding, 648
 - converting, 651
 - deleting, 652
 - editing, 649
 - ordering, 651
 - RR (see resource records)
 - RSYNC
 - password, 102
- ## S
- scopes, 355
 - adding, 356
 - defining a failover for, 360
 - defining a space for, 359
 - deleting, 363
 - duplicating (DHCPv4), 362
 - editing, 357
 - moving (DHCPv4), 363
 - options, 358
 - replicating in the IPAM, 361
 - shared network, 362
 - super-scope, 362
 - SCP
 - password, 102
 - secure dynamic update, 636
 - configuring the AD domain zone, 640
 - configuring the server, 639
 - generating a GSS-TSIG key, 638
 - preparing the AD server, 637
 - uploading a GSS-TSIG key, 638
 - security
 - firewall, 87
 - saving a backup of the appliance, 1123
 - SSH password, 101
 - tracking users' operations, 1099
 - service
 - disabling, 100
 - enabling, 100
 - starting, 100
 - stopping, 100
 - time and date configuration, 79
 - services
 - admin account, 101
 - configuration files, 108

- DNS Guardian, 105
- GSLB server, 105
- SMTP relay, 103
- SNMP server, 106
- SSL certificate, 104
- TFTP, 103
- xfer account, 102
- SFTP
 - password, 102
- Single-Server, DHCPv4, 318
- Single-Server, DHCPv6, 319
- Single-Server, DNS, 446
- smart architectures
 - DHCPv4, 305
 - DHCPv6, 305
 - DNS, 435
 - locked synchronization (DHCP), 329
 - locked synchronization (DNS), 456
- Smart Folders, 1157
- SMTP, 103
- SNMP, 106
 - configure the local agent, 107
 - metrics, 1266
 - monitoring, 1103
 - DHCP, 1274
 - DNS, 1275
 - Guardian, 1278
 - hardware, 1268
 - prerequisites, 1266
 - system, 1270
 - profile, 1101
- SOLIDserver
 - basic network configuration, 18
 - clearing SOLIDserver cache, 1120
 - first installation, 8
 - from USB, 17
 - on 3rd gen. hardware, 10
 - on 4th generation hardware, 11
 - prerequisites, 8
 - SDS-50, 8
 - front/back panel, 2
 - Linux packages
 - BIND DNS, 468
 - BIND upgrade, 476
 - ISC DHCP, 342
 - ISC upgrade, 350
 - logging in, 24
 - non-supported options, 1303
 - prerequisites, 8
 - resetting the configuration via LCD, 22
 - time and date configuration, 79
 - update macros and rules, 1120
 - upgrading, 1133
- sortlist
 - on a DNS server, 507
 - on a DNS view, 556
- spaces
 - adding, 223
 - deleting, 224
 - editing, 223
- SPF Records (see resource records)
- Split-Scope, DHCPv4, 316
- Split-Scope, DHCPv6, 321
- SPX, 995
 - AS Numbers, 159, 1014 (see aut-nums)
 - adding, 1015
 - browsing, 1014
 - deleting, 1016
 - editing, 1016
 - aut-nums (see AS numbers)
 - importing, 159
 - configuring, 996
 - changes, 999
 - classes, 996
 - connection, 996
 - rules, 996
 - importing
 - allocated networks, 156
 - assigned networks, 157
 - persons, 158
 - importing CSV data, 156
 - networks, 1006
 - adding, 1007
 - browsing, 1006
 - deleting, 1012
 - editing, 1010
 - registering changes, 1011
 - persons, 1002
 - adding, 1002
 - browsing, 1002
 - deleting, 1004
 - editing, 1003
 - registering changes, 1004
- SSH, 101
 - enable/disable, 100
 - login, 101
 - password, 101
 - password security level, 101
- SSL
 - certificate, 104, 1106
 - CSR, 1106
 - private key, 1106
- Stateless, 322
- statics, 366
 - adding, 368
 - copying (DHCPv4), 375
 - deleting, 375
 - editing, 371

- MAC address types, 1255
- options, 373
- replicating in the IPAM, 372
- statistics, 424, 713
 - (see also analytics)
 - DHCP physical server, 423
 - DNS physical server, 712
 - Guardian clients, 802
 - Guardian server, 793
- Stealth, 444
- subnetting (see networks)
- syslog, 1095
 - severity levels, 1096
- System statistics
 - appliance and services, 1097

T

- templates (IPAM), 271
 - applying templates, 278
 - creating templates, 272
 - for block-type networks, 274
 - for pools, 276
 - for spaces, 273
 - for subnet-type networks, 275
 - template classes, 271
- tracking
 - sessions, 1099
 - users, 1099
- transfer-source, 514
- transfer-source-v6, 515
- troubleshooting, 1120
 - guidelines, 1121
 - network capture, 1122
 - troubleshooting dump, 1123

U

- undo
 - restoring deleted IP addresses, 263
- uniqueness of IP address
 - preventing IP address duplication (DHCP), 401
- update
 - DHCP database
 - from the IPAM, 732
 - DNS database
 - from the DHCP, 732
 - from the IPAM, 732
 - IPAM database
 - from the DHCP, 732
 - from the DNS, 732
- update-policy, 636
- upgrade, 1133
 - prerequisites, 1133
 - troubleshooting the upgrade, 1142

- of an appliance, 1142
- of HA appliances, 1144
- upgrading an appliance, 1133
- upgrading HA appliances, 1136
- upgrading remote appliances, 1135
- use-alt-transfer-source, 515
- user, 1030
 - adding, 1031
 - changing a user password, 1033
 - configuring the user login session time, 1036
 - configuring the user password complexity, 1034
 - configuring the users connection parameters, 1034
 - connected user, 43
 - password, 44
 - settings, 43
 - editing, 1032
 - enabling/disabling, 1037
 - granting access to changes from all the users, 1100
 - groups of users, 1019
 - importing CSV data, 161
 - local user, 1030
 - redirecting user once they logged out, 1036
 - remote authentication, 1038
 - tracking user operations, 1099

V

- views, 538
 - adding, 540
 - deleting, 544
 - editing, 542
 - getting rid of all views, 545
 - order, 543
- VIF (virtual interface), 96
- VIP (virtual IP), 94
- Virtual Local Area Network
 - setting up a VLAN using a VIF, 90
- VitalQIP
 - importing DHCP data, 168
 - importing DNS data, 176
 - importing IPAM data, 164
- VLAN, 981
 - adding, 982
 - deleting, 983
 - editing, 983
 - statuses, 982
- VLAN Manager, 973
 - associating VLANs and IPAM networks, 985
 - domains (see domains (VLAN))
 - importing CSV data, 150
 - ranges (see ranges (VLAN))
 - removing the IPAM/VLAN interaction, 986
 - rights, 1023

- updating from IPAM networks, 984
- VLANs (see VLAN)
- VLSM (Variable Length Subnet Masking), 280
 - choosing the proper method, 280
 - moving IPv4 addresses across the VLSM, 293
 - network-based organization, 294
 - reparenting subnet-type networks, 296
 - space-based organization, 284
- VRF, 988
 - importing CSV data, 153
 - Virtual Routing and Forwarding
 - adding, 989
 - deleting, 990
 - editing, 990
 - importing, 154
 - VRF Route Target
 - adding, 992
 - deleting, 993
 - importing, 155
- VRRP (Virtual Router Redundancy Protocol), 93
- VXLAN (see VLAN)

W

- Workflow, 910
 - customizing the requests administration, 930
 - adding statuses, 934
 - best practices, 935
 - editing the email notification, 933
 - editing the statuses, 931
 - executing requests, 925
 - using classes, 926
 - using the execute option, 925
 - incoming requests, 920
 - accepting requests, 922
 - archiving requests, 923
 - default request administration, 921
 - deleting requests, 923
 - finishing requests, 923
 - handling requests, 922
 - managing requests content, 920
 - rejecting requests, 923
 - outgoing requests, 912
 - adding creation requests, 913
 - adding deletion requests, 915
 - adding edition requests, 914
 - canceling requests, 919
 - editing requests, 917
 - user access to classes, 911
- WSDL files, 1115
 - adding, 1115
 - deleting, 1118
 - downloading, 1119
 - editing and dumping, 1116

X

- X.509
 - HTTPS certificate, 1106

Z

- zones
 - adding
 - delegation-only, 576
 - forward, 569
 - hint, 575
 - master, 563
 - slave, 567
 - stub, 572
 - adding or removing an NS record, 586
 - administrator rights on, 595
 - classless in-addr.arpa delegation, 597
 - converting, 585
 - delegation, 596
 - deleting, 595
 - disabling/enabling, 594
 - duplicating zones, 586
 - editing
 - delegation-only, 584
 - forward, 581
 - hint, 583
 - master, 579
 - slave, 580
 - stub, 582
 - forcing a full synchronization, 578
 - forcing commands, 593
 - notification, 593
 - refresh, 593
 - retransfer, 594
 - hosting active directory domain zones, 578
 - importing a VitalQIP archive file, 177
 - migrating to another server, 586
 - moving zones, 587
 - security, 602
 - setting authorizations, 587
 - allow-query, 588
 - allow-transfer, 588
 - allow-update, 589
 - setting forwarders, 592
 - setting master servers, 592
 - setting properties, 587
 - setting space, 587
 - statuses, 562
 - synchronizing the manager with, 578