

【CSC公开课】第十一期

# iSCSI Boot UCS-B with Cisco VIC with Network Appliance Filers



Planning and Deployment

# iSCSI Boot

## Coverage and Assumptions

- This document covers booting a UCS B-Series blade into VMware ESXi from a Network Appliance filer over iSCSI.
  - The concepts for the UCS configuration are applicable to other iSCSI capable storage arrays.
- Sections
  - Overview
  - UCS Profile Definition
  - NetApp Configuration
  - Troubleshooting

# iSCSI Boot

## Disclaimers

- UCS Manager 2.0(1m) or higher is required.
- The functionality is supported for the UCS Virtual Interface Card (VIC) and Broadcom adapters.
  - This document covers the Cisco VIC.
- Consult the UCS Interoperability Matrix to validate the storage target firmware versions.

# iSCSI Boot

## Pre-Requisites

- UCS Manager 2.0(1m) or higher
- VMware 4.1 or higher
  - (Recommended) Use VMware image that has the Cisco drivers preloaded.  
<https://my.vmware.com/web/vmware/downloads>  
Search for “**cisco customized**”
- Cisco Virtual Interface Card (VIC) on UCS B-Series blade
  - (2) IP Addresses for iSCSI Initiators/server
- Network Appliance filer licensed for iSCSI
- Administrative access to UCS and Filer

# iSCSI Boot UCS-B with Cisco VIC

## Overview



Planning and Deployment


# iSCSI Boot

## Overview

- iSCSI Boot provides a means to host the server's operating system on centralized, highly available storage without requiring a Fibre Channel infrastructure.
- Remote booting UCS servers provides additional mobility for a service profile as the operating system no longer physically resides on the server's local drives.
- The goal of this document is to provide a comprehensive description for design, deployment and troubleshooting iSCSI boot in UCS Manager.

# iSCSI Boot

## Mezzanine Card



- The DCE interfaces go from the physical blade to the UCS IO Module/Fabric Extender. The vNICs face the operating system and Option ROM.
- Each iSCSI initiator requires an IP address.
  - They can be assigned Statically (Recommended) or Dynamically (DHCP)

# iSCSI Boot

## Design Options – Path Optimizations

- There are design options that can be applied to the configuration and should be **applied to the entire path** from the initiator (server) to UCS to upstream switch(es) to the target (storage array).
- Jumbo Frames (Recommended)
  - Jumbo Frames of 9000 bytes are often used to improve the efficiency loading the operating system quickly.
- Traditional Ethernet (Recommended)
  - Traditional Ethernet is inherently fair in terms of giving each consumer a relatively equal slice of the network pipe. In a 10GbE pipe, there is typically more room so the fairness is more pronounced.
  - If iSCSI is being used only as the storage protocol for booting, the default traditional Ethernet settings should be acceptable.
- Lossless Ethernet (Optional)
  - The UCS and Nexus switches support Lossless Ethernet which is used by Fibre Channel over Ethernet to ensure no frames are dropped.
  - Lossless Ethernet can be considered if there are a large number of hosts using iSCSI boot and there is a desire to insure against boot storms or if iSCSI is being used for both boot and data.



# iSCSI Boot

## Design Options – Path Selection

- The UCS VIC supports a feature called Fabric Failover wherein the VIC can switch the traffic on a vNIC from one Fabric Interconnect to the other in the event there is a loss of connectivity to the vNIC's parent path. This functionality is performed in hardware.
- The Cisco recommendation is to **not** use this functionality for Fibre Channel, FCoE and iSCSI but rather to rely on multi-path input/output (MPIO) drivers that come with the operating system or from third party vendors to make the storage path decision.
- Cisco does not officially test fabric failover with these storage protocols.

# iSCSI Boot

## Theory of Operation

- UCS will program the VIC to present two iSCSI initiators (A/B) path.
- The initiators are programmed with the initiator and destination IQNs.
- The boot policy is set to boot from the iSCSI vNICs.
- When the server boots, the VIC's software will log into the storage array on both paths.
- The primary will be used to load the operating system (or the secondary if the primary fails to login).
- Once the operating system is loaded, the VIC's boot loader transfers control to the operating system. At this point the operating system manages which vNIC is being used and path failover.
- The iSCSI vNIC that is not being used will eventually time out and on the filer, only one IQN will be seen for the given host.

# iSCSI Boot UCS-B with Cisco VIC Upstream Switch

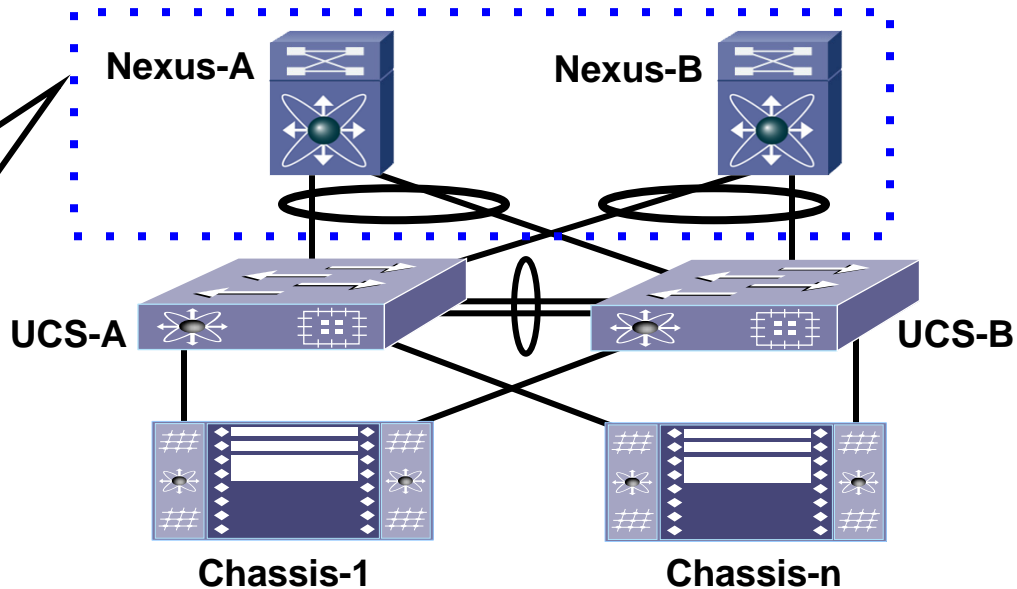


Planning and Deployment

# iSCSI Boot – Upstream Switch

## Jumbo Frames

```
policy-map type network-qos class-jumbo
  class type network-qos class-default
    mtu 9000
  system qos
  service-policy type network-qos class-
jumbo
```



- Configure the UCS' s upstream switches to pass jumbo frames.
  - The example above shows a Nexus switch running NX-OS.
- Save the configuration on each switch.
- (Recommended) Use Virtual Port Channels from the UCS to the upstream switches if possible to provide more resiliency for the storage connectivity.

# iSCSI Boot – Upstream Switch

## Jumbo Frames – Verify

```
nexus5548-11a# show policy-map type network-qos class-jumbo
```

```
Type network-qos policy-maps
```

```
=====
```

```
policy-map type network-qos class-jumbo  
class type network-qos class-default
```

```
mtu 9000
```

```
multicast-optimize
```

- Verify the jumbo frames are enabled.
- For a Nexus running FCoE, there are jumbo frames for FCoE and all other traffic (Default).

# iSCSI Boot UCS-B with Cisco VIC UCS Manager

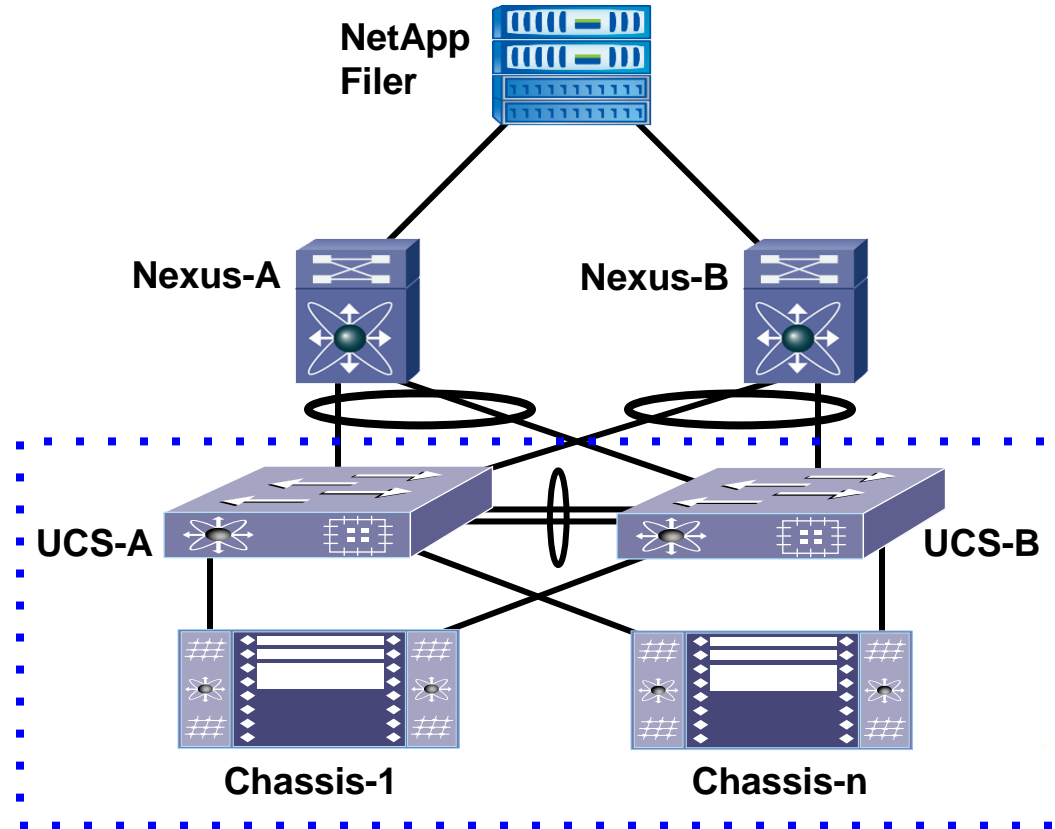


Planning and Deployment

# iSCSI Boot – UCS

## Overview

- The UCS cluster must be configured to allow Jumbo Frames to pass through to the storage.
- (Recommended) A vNIC template should be used to ensure the settings are consistently deployed.
- A boot policy is defined for iSCSI boot which includes the definition of the target storage array and LUN.



# iSCSI Boot – UCS

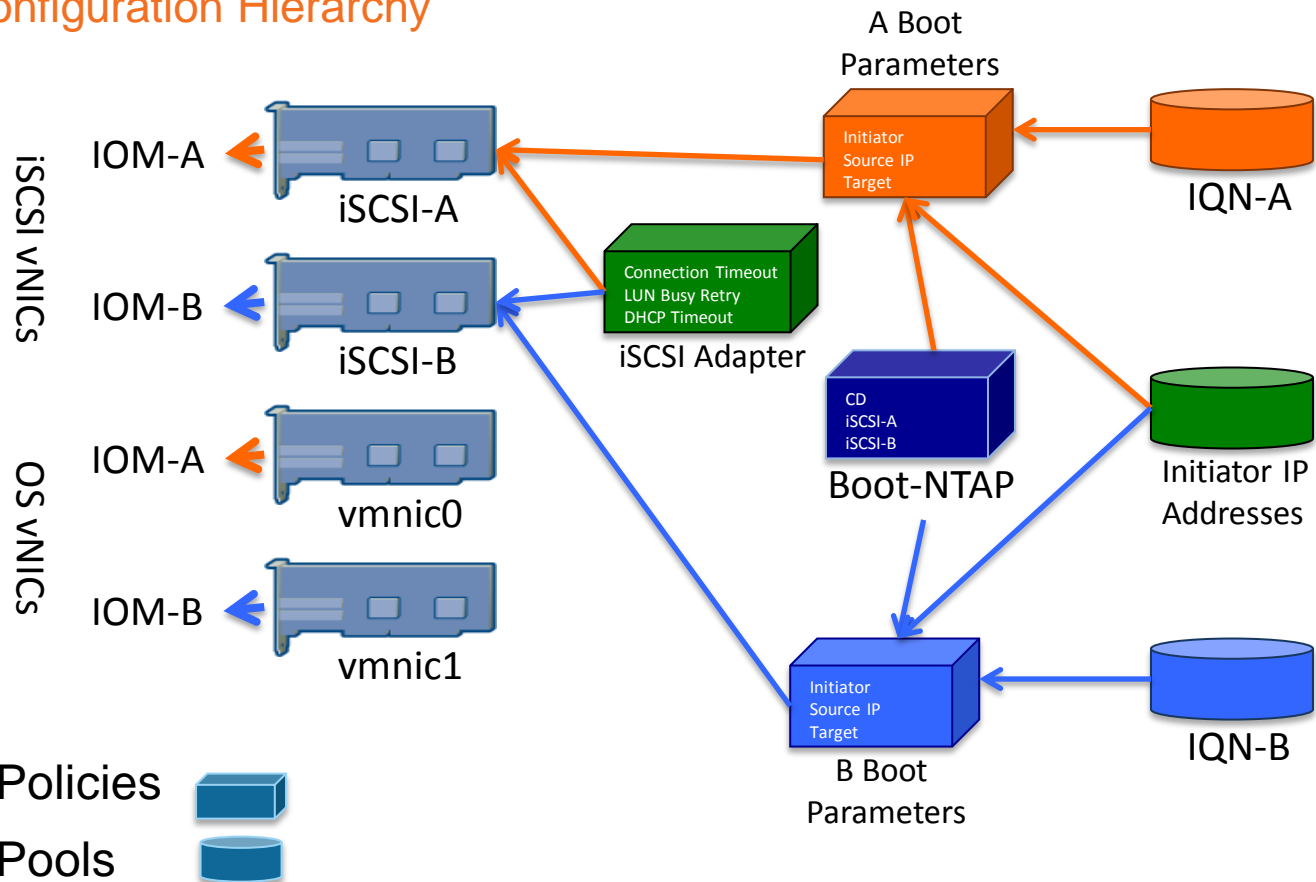
## Configuration Overview

- **Pools** are created for the Initiator IP Addresses and IQNs.
  - The Initiator Pool is global across all organizations. The initiators support a default gateway in the event the Initiator Pool and the target storage array are not on the same subnet.
- **Policies** are created to govern connectivity, mapping the vNIC to the iSCSI initiator, the iSCSI target, define boot order and, optionally, security to access the LUN.
- A **Service Profile** leverages the Pools and Policies to establish the personality for the server.
- (Optional) A Service Profile can be used to create a **Service Profile Template** which can be used to instantiate multiple instances of Service Profiles using a common base definition.



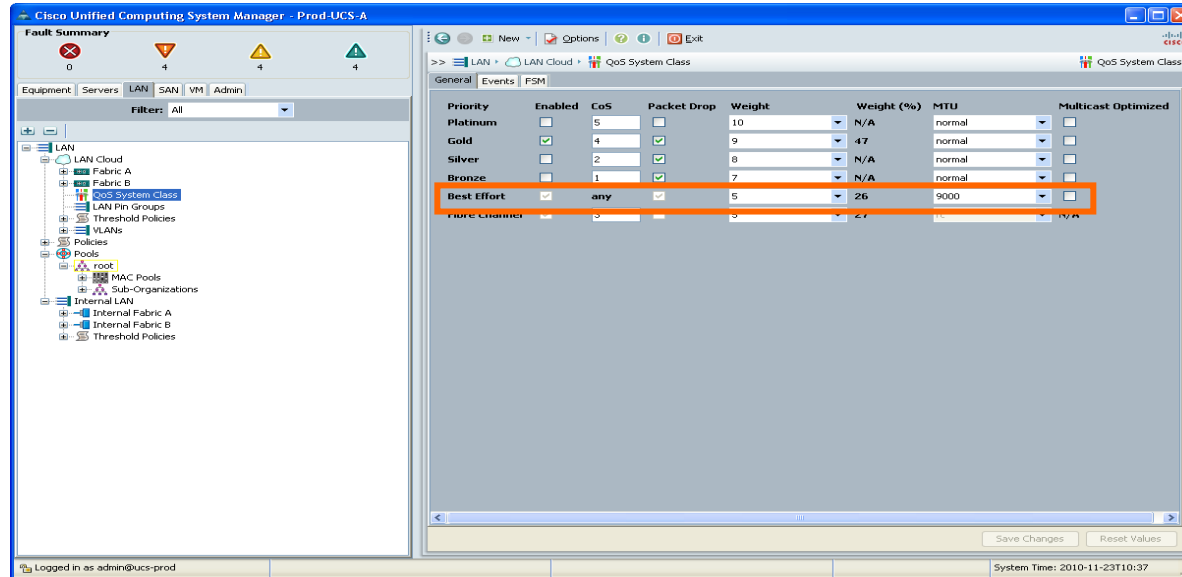
# iSCSI Boot – UCS

## Configuration Hierarchy



# iSCSI Boot – UCS

## Jumbo Frames – System Wide



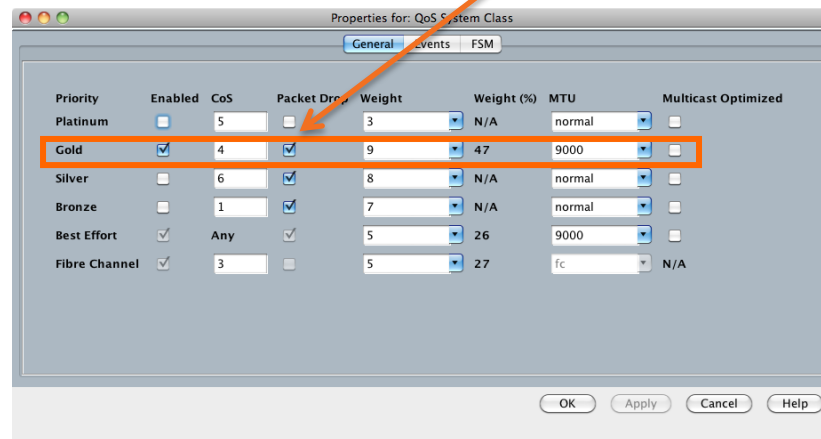
- Enable Jumbo frames on the UCS cluster as a whole.
  - LAN → LAN Cloud → QoS System Class

# iSCSI Boot – UCS

## Jumbo Frames – System Wide – iSCSI at Scale (Optional)

- If the iSCSI boot will be used for a large scale deployment, to insulate against boot storms, finer grained Quality of Service may be of interest.
- Define a queue for the IP Storage Traffic.
- Set the MTU to 9000.
- Set the queue to be non-drop by unchecking **Packet Drop**.
- If FCoE will not be used, consider lowering the weight for FCoE and transferring it to the Priority queue being used for iSCSI.

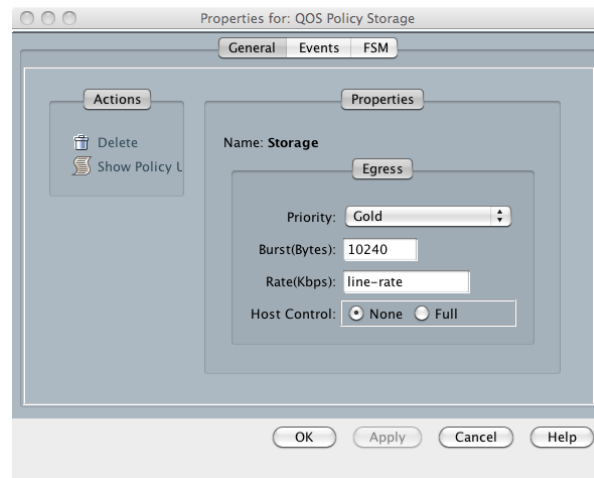
Non-Drop



# iSCSI Boot – UCS

## QoS Policy – iSCSI at Scale (Optional)

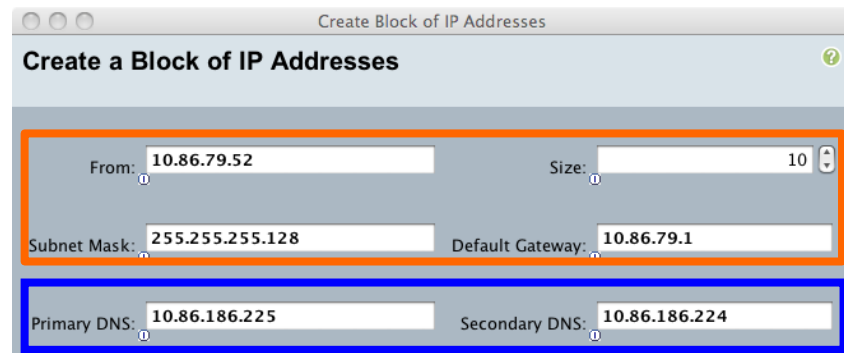
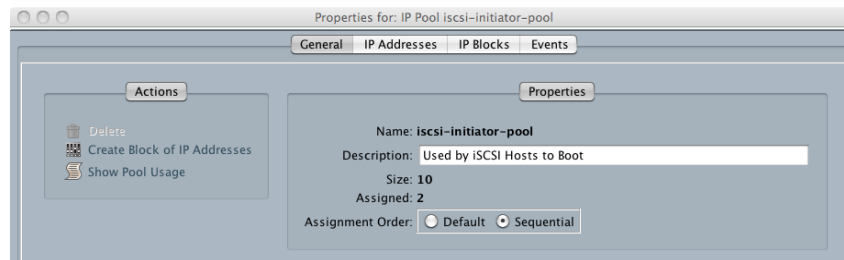
- In **LAN** → **Policies** → **<Org>** → **QoS Policies**, create a new policy called **Storage**.
- Set the Priority to the queue chosen in **QoS System Class**
  - **Gold** in this example.
- Be sure to set the upstream switch up to also support this non-drop traffic class.



# iSCSI Boot – UCS

## LAN Policies – iSCSI Address Pool

- In LAN → Pools → root → IP Pool **iscsi-initiator-pool** create a new block of IP addresses.
  - The IP addresses can be on the same subnet as the target (recommended) or on a routed network the server can access.
  - Budget two IP addresses per expected host. The pool can be expanded later.
  - (Optional) Set the Pool Assignment order to **Sequential**.
- (Recommended) Add the addresses into DNS.

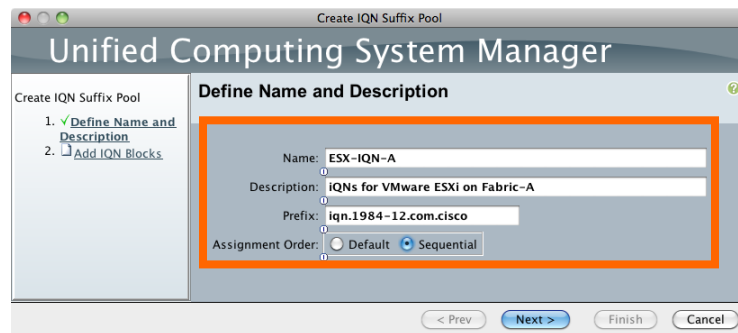


❗ Be sure to record the IP addresses in DNS to avoid someone else using them

# iSCSI Boot – UCS

## SAN Policies – iSCSI IQN Pool A

- In **SAN** → **Pools** → **<Org>** → **IQN Pools** create a new block of IQNs for the A and B path.
- Give the pool a name.
  - Filers expect a valid prefix. Use **iqn.1984-12.com.cisco**
  - Specify the IQN's prefix.
  - (Recommended) Set the Assignment Order to **Sequential** as this will make it easier to determine which address will be consumed next.
  - Press **Next**



# iSCSI Boot – UCS

## SAN Policies – iSCSI IQN Pool A

- Select Add and specify the suffix
  - Operating System, Application, Organization, etc.
- The **From** field specifies the start of the range.
  - In this case we have vmware01-34 so we are starting with 35

Create a Block of IQN Suffixes

Create a Block of IQN Suffixes

Suffix:

From:

Size:

OK Cancel

Create IQN Suffix Pool

Unified Computing System Manager

Create IQN Suffix Pool

1. [Define Name and Description](#)
2. [Add IQN Blocks](#)

Add IQN Blocks

Name	From	To
vmware-prod-a:35 - vmv35	35	42


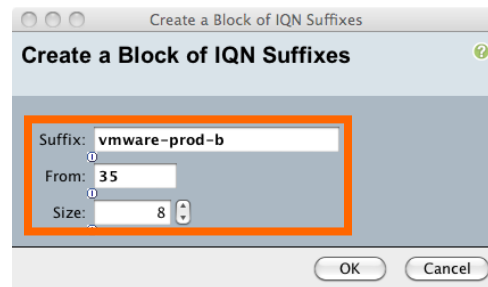
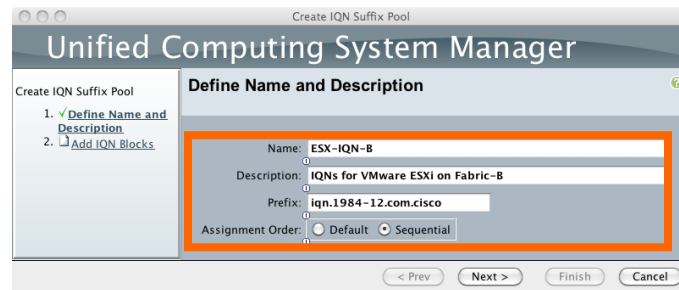
Add Delete

< Prev Next > Finish Cancel

# iSCSI Boot – UCS

## SAN Policies – iSCSI IQN Pool B

- Repeat, creating an IQN Pool for the operating system on the “-B” path.
- Make sure the pool size and assignment order is the same as the A Pool.
  - This will ensure the sequence numbers stay in sync as the pools are consumed.
- ⓘ It is not a requirement to have two pools but it helps in troubleshooting when looking at the filer to see which path is logged in.





# iSCSI Boot – UCS

## LAN Policies – vNIC Service Template

- In **LAN → Policies → <Org> → vNIC Templates**, define vNIC templates for the host to use for VMware for each path (A/B).
  - This will help with debugging
- Specify the VLAN's required.
- Specify the iSCSI Boot VLAN and make it the Native VLAN.
- Set the MTU to 9000 bytes
- Specify a MAC Pool.
- Do not set Enable Failover.

# iSCSI Boot – UCS

## LAN Policies – vNIC Template

**Create vNIC Template**

Name: ISCSI-Boot-A

Description:

Fabric ID:  Fabric A  Fabric B  Enable Failover

Target:  Adapter

Warning: If VM is selected, a port profile by the same name will be created. If a port profile of the same name exists, and updating template is selected, it will be overwritten.

Template Type:  Initial Template  Updating Template

Select	Name	Native VLAN
<input type="checkbox"/>	Pod-402	<input type="checkbox"/>
<input type="checkbox"/>	Pod-403	<input type="checkbox"/>
<input type="checkbox"/>	Pod-404	<input type="checkbox"/>
<input type="checkbox"/>	Pod-405	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Servers	<input checked="" type="checkbox"/>

MTU: 9000

MAC Pool: Demonstration-MAC

QoS Policy: <not set>

Network Control Policy: CDP

Pin Group: <not set>

Stats Threshold Policy: default

OK Cancel

Different Fabrics  
① Fabric Failover is disabled

**Create vNIC Template**

Name: ISCSI-Boot-B

Description:

Fabric ID:  Fabric A  Fabric B  Enable Failover

Target:  Adapter

Warning: If VM is selected, a port profile by the same name will be created. If a port profile of the same name exists, and updating template is selected, it will be overwritten.

Template Type:  Initial Template  Updating Template

Select	Name	Native VLAN
<input type="checkbox"/>	Pod-402	<input type="checkbox"/>
<input type="checkbox"/>	Pod-403	<input type="checkbox"/>
<input type="checkbox"/>	Pod-404	<input type="checkbox"/>
<input type="checkbox"/>	Pod-405	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Servers	<input checked="" type="checkbox"/>

MTU: 1500

MAC Pool: Demonstration-MAC

QoS Policy: <not set>

Network Control Policy: CDP

Pin Group: <not set>

Stats Threshold Policy: default

OK Cancel

Native VLAN

MTU

- Create the templates for the A/B fabrics (LAN → Policies → <Org> → vNIC Template).

# iSCSI Boot – UCS

## Server Policies – Adapter Policy

Create iSCSI Adapter Policy

**Create iSCSI Adapter Policy**

Name: VIC-ESXI

Connection Timeout: 0 [0-255]

LUN Busy Retry Count: 0 [0-60]

DHCP Timeout: 60 [60-300]

Enable TCP Timestamp:

HBA Mode:

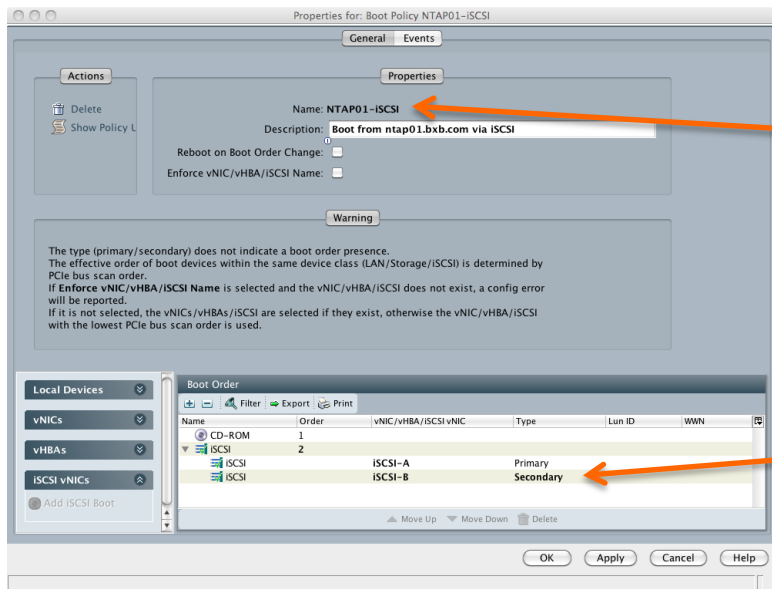
Boot To Target:

OK Cancel

- In **Servers** → **Policies** → <Org> → **Adapter Policies** → **iSCSI Adapter Policies**, select **Create**.
- Give the policy a name.
- Leave **Boot to Target** unchecked (this used by Broadcom).

# iSCSI Boot – UCS

## Server Policies – Boot



① Use a naming convention for the boot policy (e.g.- <Array>- <Protocol>)

① The vNICs will be created later

- Create a boot policy in **Servers → Policies → <Org> → Boot Policies**.
  - Specify CD-ROM first to boot the installation media.
  - Expand **iSCSI vNICs** and add two vNICs.
- Note that there is no option to specify the target like with a vHBA. This is done when the boot profile is consumed.
  - The target definition will be manually configured at profile creation time.

# iSCSI Boot – UCS

## Service Profile – Server – Name

- Create a service profile (**Servers** → **Service Profiles** → **<Org>** → **Create Service Profile (expert)**).
- Select a UUID.
- (Optional) Specify a description.

Unified Computing System Manager

Create Service Profile (expert)

**Identify Service Profile**

You must enter a name for the service profile. You can also specify how a UUID will be assigned to this profile and enter a description of the profile.

Name:

The service profile will be created in the following organization. Its name must be unique within this organization.

Where: **org-root/org-Demonstration**

Specify how the UUID will be assigned to the server associated with this service profile.

UUID Assignment:

Create UUID Suffix Pool

The UUID will be assigned from the selected pool.  
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev Next > Finish Cancel

# iSCSI Boot – UCS

## Service Profile – Server – Add OS Facing vNICs

Unified Computing System Manager

Create Service Profile (expert)

1. Identify Service Profile  
2. **Networking**  
3. Storage  
4. zoning  
5. vNIC/iHRA Placement  
6. Server Boot Order  
7. Maintenance Policy  
8. Server Assignment  
9. Operational Policies

**Networking**  
Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) [Create Dynamic vNIC Connection Policy]

How would you like to configure LAN connectivity?  Simple  **Expert**  No vNICs  Hardware Inherited  Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
vNIC vmnic0	Derived	derived	
vNIC vmnic1	Derived	derived	

Click **Add** to specify one or more iSCSI vNICs that the server should use.

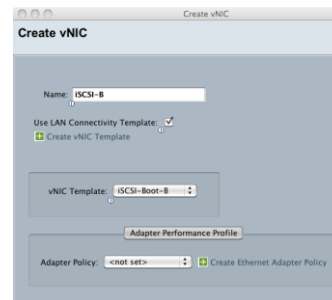
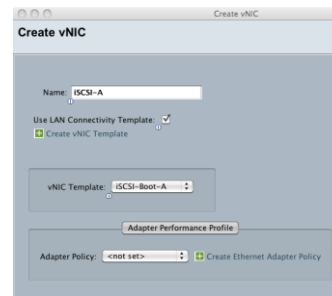
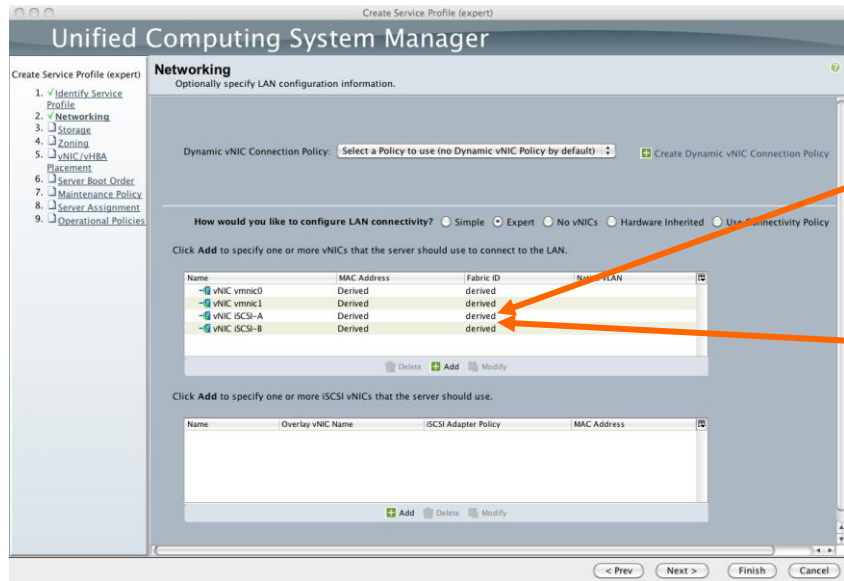
Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
------	-------------------	----------------------	-------------


< Prev Next > Finish Cancel

- Select the **Expert** option.
- Create the Ethernet vNICs the operating system will use.

# iSCSI Boot – UCS

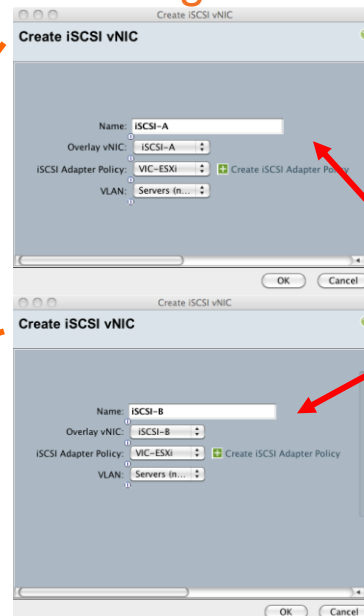
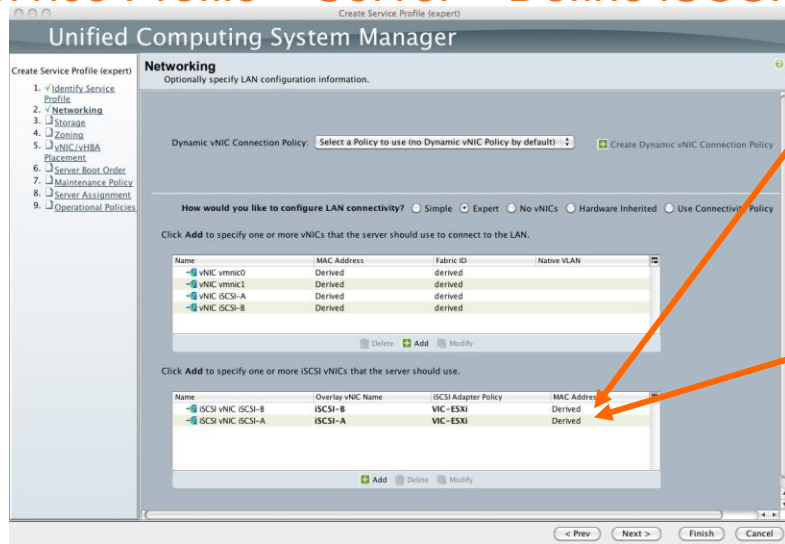
## Service Profile – Server – Add iSCSI vNICs



- Select the LAN Expert and use the Templates that were previously defined.
- Repeat for the B Fabric NIC.
-  If manually defining the NICs remember to set the MTU to 9000 bytes.

# iSCSI Boot – UCS

## Service Profile – Server – Define iSCSI vNIC Settings



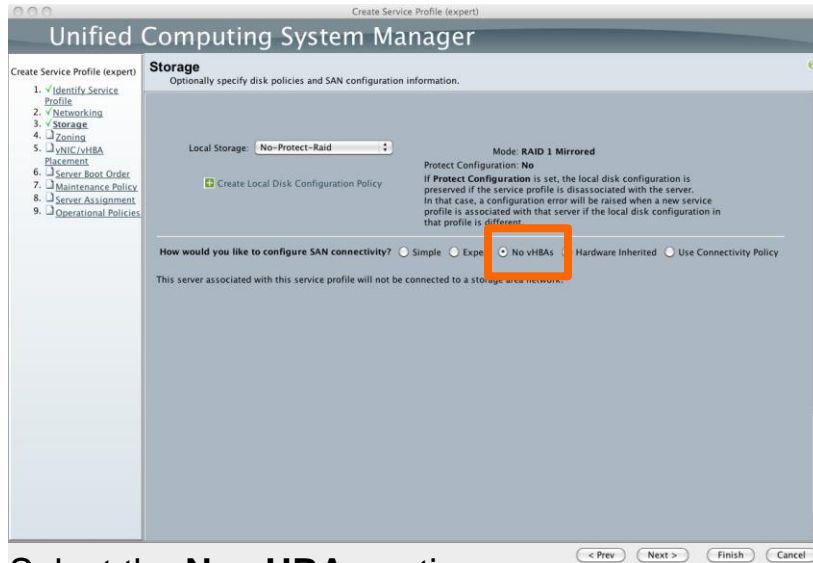
❗ Ensure the Name matches the iSCSI vNIC name in the boot policy.

- Expand the iSCSI vNIC section and connect the iSCSI vNICs to their policy/Native VLAN.
- Create a link for each vNIC specifying the overlay iSCSI vNIC adapter
  - ❗ Do not use the Default VLAN. Use the Native VLAN defined for the parent NIC.
  - ❗ Do not set the MAC address. The Overlay iSCSI vNIC uses the parent vNIC when using the Cisco VIC.



# iSCSI Boot – UCS

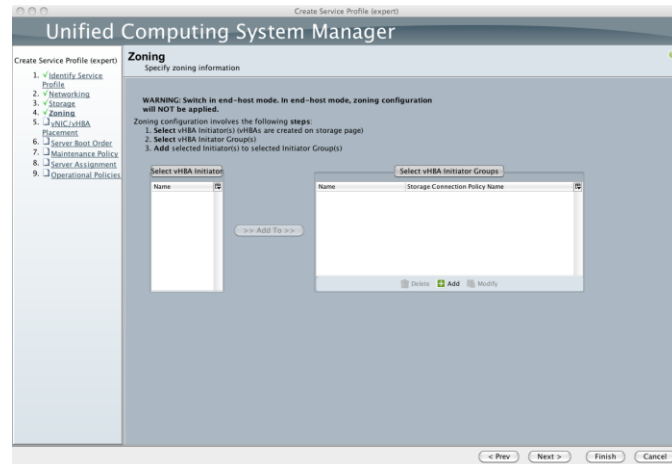
## Service Profile – Server – Local Storage



- Select the **No vHBAs** option.
  - The storage will come through the vNIC.

# iSCSI Boot – UCS

## Service Profile – Server – Zoning

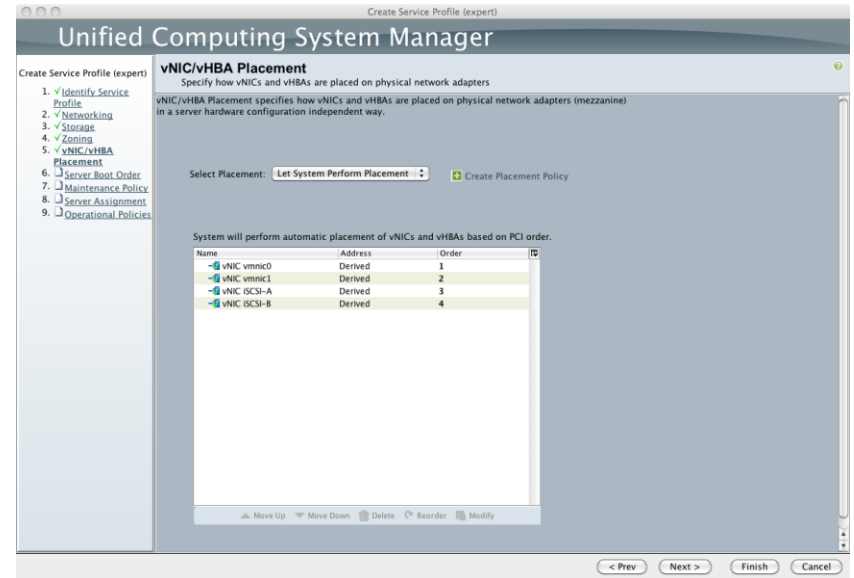


- Skip the **Zoning** page as fibre channel is not being used.

# iSCSI Boot – UCS

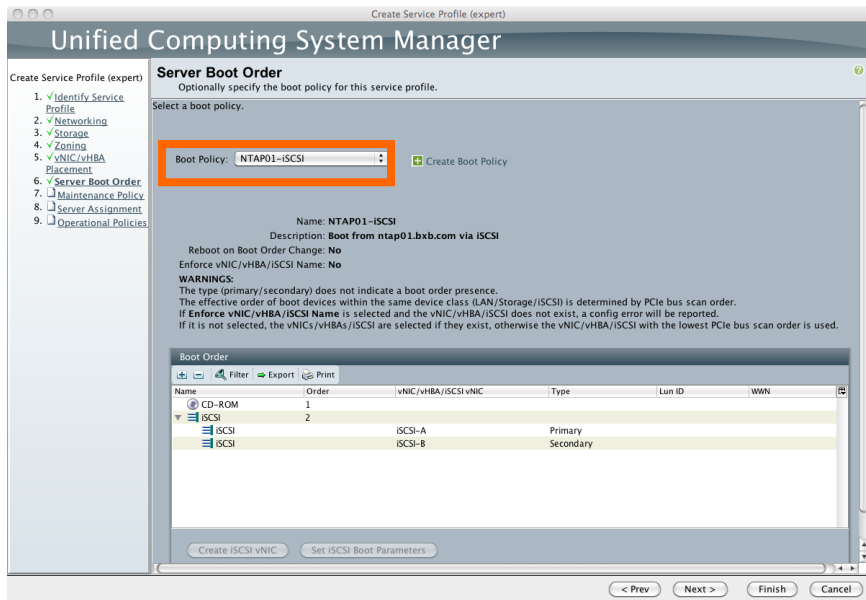
## Service Profile – Server – vNIC Placement

- Leave the default placement.



# iSCSI Boot – UCS

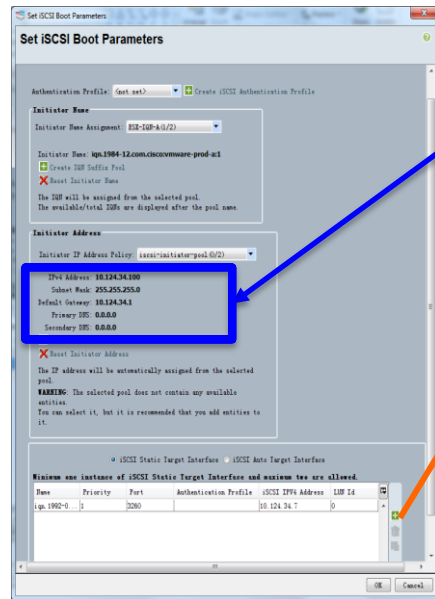
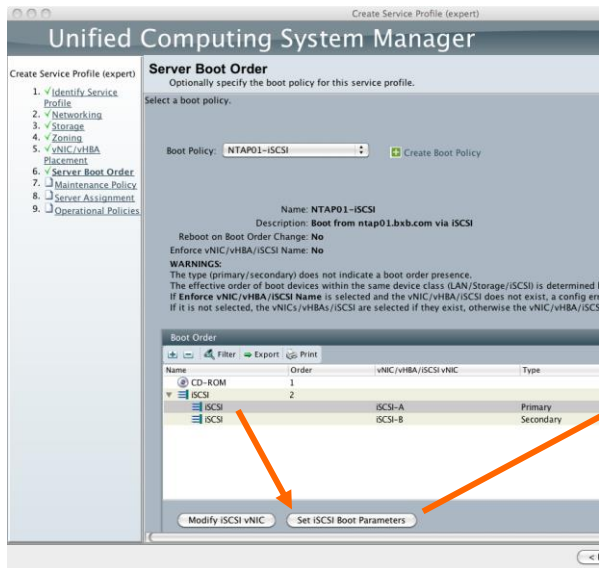
## Service Profile – Server – Use iSCSI Boot Policy



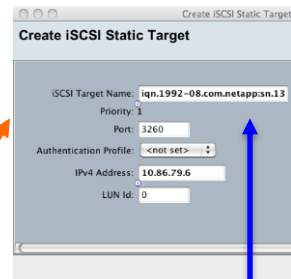
- Select the previously created boot policy.

# iSCSI Boot – UCS

## Service Profile – Server – Define Boot Target – A



The IP address information will be filled in after the profile is created.



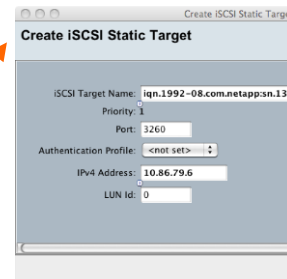
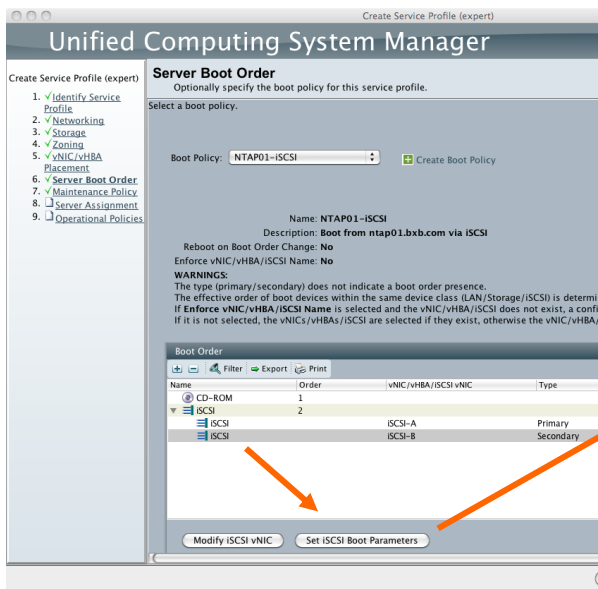
```
ntap01> iscsi nodename  
iSCSI target nodename: iqn.1992-08.com.netapp:sn.135077524
```

From the storage array

- Select the A iSCSI vNIC and **Set iSCSI Boot Properties**
  - If the button stays grayed out the vNIC Name in the Boot policy probably does not match the Overlay iSCSI vNIC name in the Network step.
- Specify the A IQN Suffix pool, the Initiator's IP address
- Create a target definition that points to the storage array's iSCSI subsystem.

# iSCSI Boot – UCS

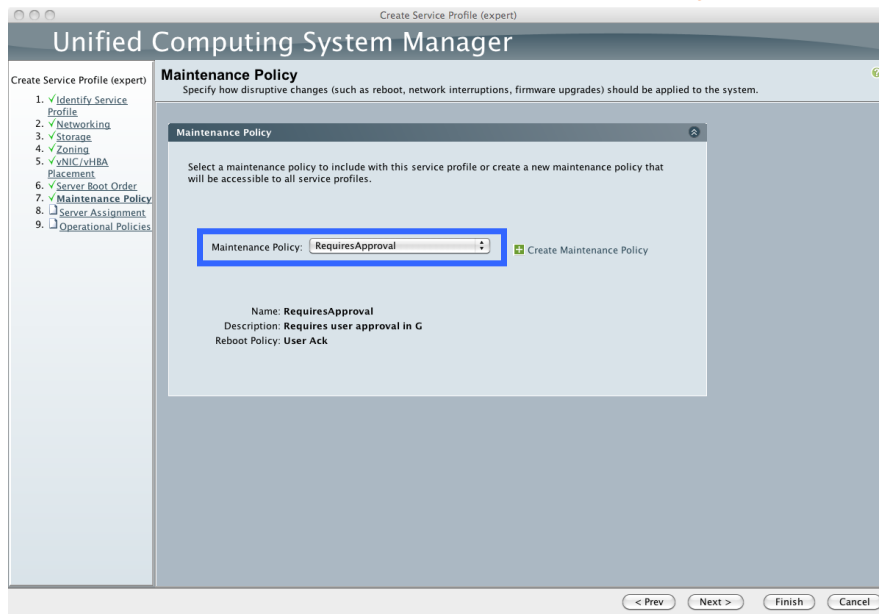
## Service Profile – Server – Define Boot Target - B



- Select the B iSCSI vNIC and **Set iSCSI Boot Properties**
- Specify B IQN Suffix pool, the Initiator's IP address
- Create a target definition that points to the same storage array.

# iSCSI Boot – UCS

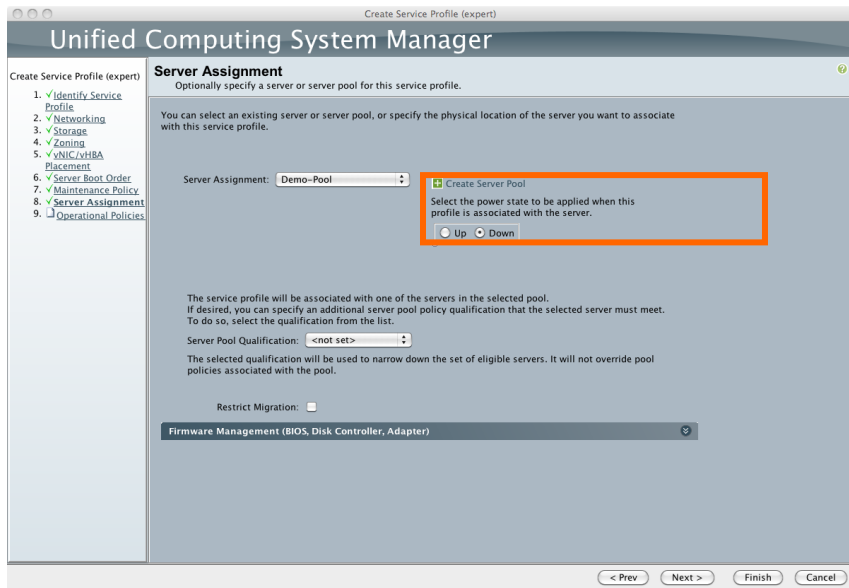
## Service Profile – Server – Maintenance Policy



- (Optional) Specify a maintenance policy.
  - While creating a first profile, this can be left at defaults.
  - After the profile is validated and is going to be made into a Service Profile Template, it is recommended to set a User Ack maintenance policy.

# iSCSI Boot – UCS

## Service Profile – Server – Hardware Assignment

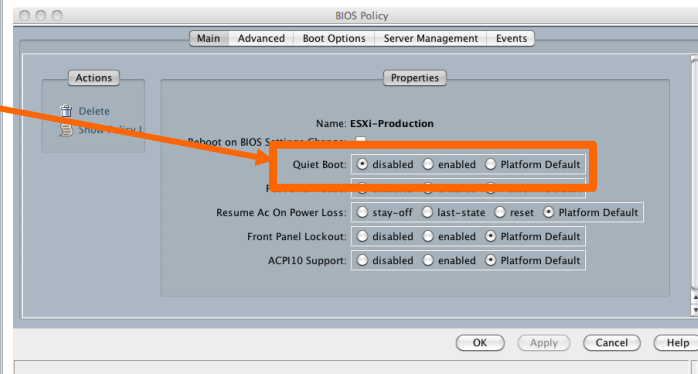
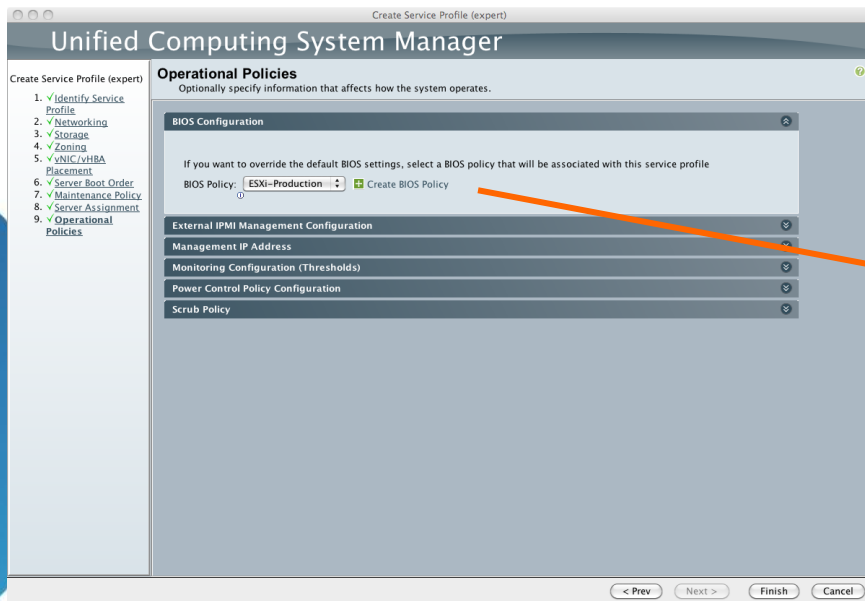


- Choose a server or a server pool.
  - (Recommended) Use a server pool.
- (Optional) Select the **Down** radio button to give time to setup the KVM and/or storage.



# iSCSI Boot – UCS

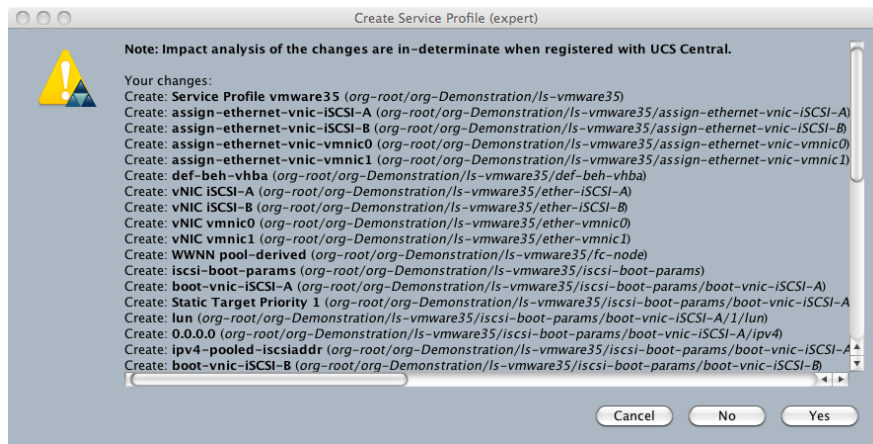
## Service Profile – Server – Disable Quiet Boot



- (Optional) Select a BIOS policy that disables quiet boot so the Option ROM messages can be seen.

# iSCSI Boot – UCS

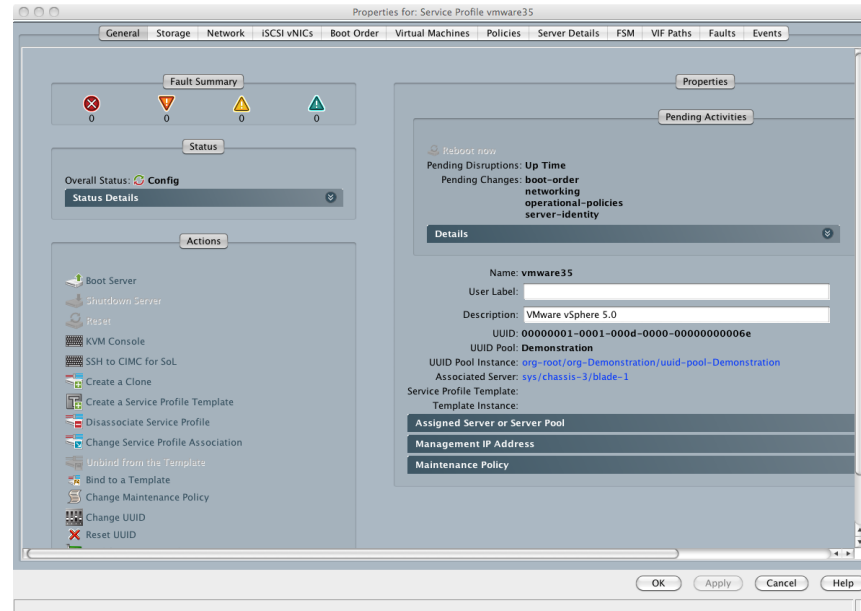
## Service Profile – Server – Confirm Association



- Select **Yes** to confirm the provisioning of the server.

# iSCSI Boot – UCS

## Service Profile – Server



- While the service profile is deploying, configure the storage.

# iSCSI Boot UCS-B with Cisco VIC Storage Setup




Planning and Deployment

# iSCSI Boot – NetApp

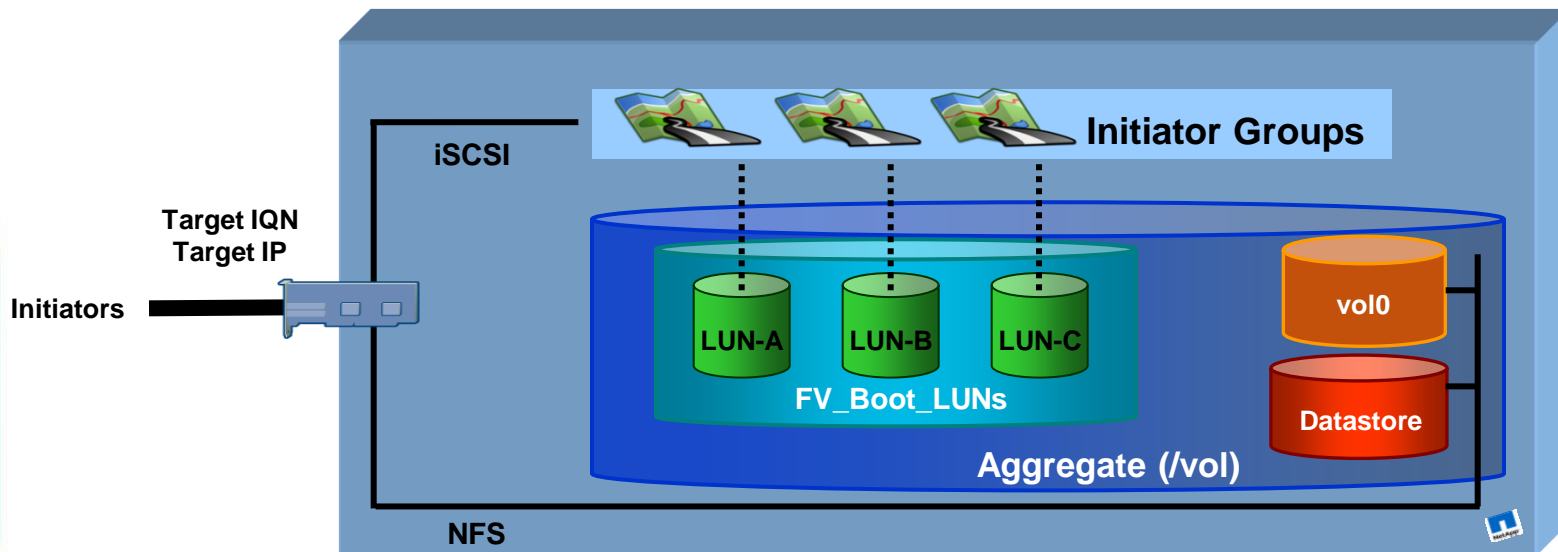
## Overview

- Configure the Network Appliance filer to present an iSCSI boot LUN to the host.
- This section will use OnTAP CLI commands



# iSCSI Boot – NetApp


## Configuration Hierarchy



- Aggregates group physical disks.
- They contain volumes which contain like data typically.
- The Boot LUNs are mapped to initiator groups which specify the IQN that is allowed to use them and which LUN number to present to the validated initiator.

# iSCSI Boot – NetApp

## Configuration Hierarchy – Details



- The initiator group (iGroup) defines which initiators will use a common operating system
- The LUN OS type and iGroup must be the same type.
- To preserve a 1:1 mapping with LUN0 as the boot LUN, an iGroup is created for each host.
  - The initiators are the primary and secondary NICs on that host.

# iSCSI Boot – NetApp

## Validate Environment

```
ntap01> aggr status
      Aggr State      Status      Options
  aggr0 online      raid_dp, aggr      root

ntap01> iscsi status
iSCSI service is running

ntap01> iscsi interface show
Interface e0a disabled
Interface e0b disabled
Interface e1 enabled
ntap01>
```

- Use the **aggr status** command to locate the storage aggregate to use.
- The **iscsi status** and **iscsi interface show** commands will verify iSCSI is enabled and available on an interface.



# iSCSI Boot – NetApp

## Record Target Settings

```
ntap01> ifconfig e1
e1: flags=0x3d48867<UP, BROADCAST, RUNNING, MULTICAST, TCPCKSUM> mtu 9000
    inet 10.86.79.6 netmask 0xfffff80 broadcast 10.86.79.127
    ether 00:07:43:06:25:e3 (auto-10g_sr-fd-up) flowcontrol full

ntap01> iscsi nodename
iSCSI target nodename: iqn.1992-08.com.netapp:sn.135077524
```

- Use **ifconfig <Interface\_from\_iSCSI\_Interface>** to obtain the IP address.
- The **iscsi nodename** will show this filer's target IQN.
- Record both values as they will be needed when the initiator is configured.

# iSCSI Boot – NetApp

## Create the LUN Volume

Volume for Boot LUNs      Aggregate      Volume Size

```
ntap01> vol create FV_Demo_LUNs -s volume aggr0 30g  
Creation of volume 'FV_Demo_LUNs' with size 30g on containing aggregate  
'aggr0' has completed.
```

- Create a volume to host the boot LUNs if one does not already exist.
- Specify the volume name FV\_Demo\_LUNs above.
- Specify the aggregate to host the Volume (e.g.- aggr0)
- And finally specify the total size.

# iSCSI Boot – NetApp

## (Optional) Extend the LUN Volume

Volume for Boot LUNs      Grow Volume

```
ntap01> vol size FV_Demo_LUNs +20g  
vol size: Flexible volume 'FV_Demo_LUNs' size set to 50g.
```

- If a volume already exists for boot LUNs, extend it by the size needed to hold the new boot LUNs that are being provisioned.

# iSCSI Boot – NetApp

## Create the Boot LUN(s)

Boot LUN Size      LUN Consumer      Boot LUN Path

```
ntap01> lun create -s 10g -t linux /vol/FV_Demo_LUNs/vmware35
```

- Within the volume, create a boot LUN of size 10GB
- Create the LUN with a type of **linux**.
  - The “vmware” type would be used for VMware data stores.
  - VMware ESX is a Linux operating system
- (Recommended) Set the LUN name to be the same as the service profile name.

# iSCSI Boot – NetApp

## Create the Initiator Group

```
ntap01> igroup create -i -t linux vmware35
```

```
ntap01> igroup add vmware35 iqn.1984-12.com.cisco:vmware-prod-a:35
```

```
ntap01> igroup add vmware35 iqn.1984-12.com.cisco:vmware-prod-b:35
```

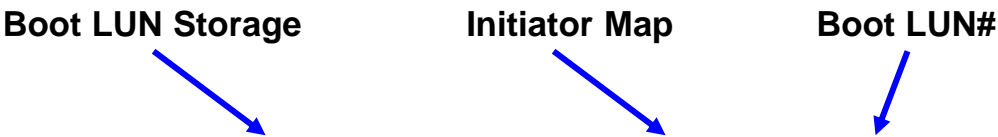
- Create an initiator group to allow the mapping of the initiator's IQN to a destination LUN.
  - The “-i” is for iSCSI.
  - Create the map of operating system type linux.
    - The “vmware” type would be used for data stores.
    - Linux will be used for the boot image.
  - (Recommended) Name the initiator group after the host/service profile.
  - Tie the map to the IQN of the initiator.
- Add the IQN to the initiator group.
  - The IQN's initiator can be any string after the vendor identifier (...com.cisco:), but the value **MUST** be consistent with the IQN suffix defined in UCS.
  - (Recommended) Use <OS>-[a|b]:<id>.

# iSCSI Boot – NetApp

## Map the LUN

Boot LUN Storage                      Initiator Map                      Boot LUN#

```
ntap01> lun map /vol/FV_Demo_LUNs/vmware35 vmware35 0
Tue Nov 23 18:45:40 EST [ntap01: lun.map:info]: LUN /vol/FV_Demo_LUNs/vmware35
was mapped to initiator group vmware35=0
```



- Tie the LUN, the Map and the LUN ID together.
  - This makes the LUN visible to the initiator.
  - The LUN must match the LUN number defined on UCS in **<Service Profile>** → **Boot Order** → **iSCSI** → **<iSCSI vNIC>** → **Set iSCSI Boot Parameters** → **iSCSI Static Target Interface** → **LUN id.**

# iSCSI Boot – NetApp

## Verify the LUN

```
ntap01> lun show -v /vol/FV_Demo_LUNs/vmware35
/vol/FV_Demo_LUNs/vmware31    10g (10737418240) (r/w, online, mapped)
  Comment: "UCS-Prod / Demonstration Org"
  Serial#: P3ecK4/nvXC3
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: linux
  Maps: vmware35=0
```

- (Optional) View the LUN mapping.
  - Verify LUN is read/write, online and mapped
  - Verify the type is Linux

# iSCSI Boot UCS-B with Cisco VIC Booting



Planning and Deployment




# iSCSI Boot – Booting

## Overview

- Installing the operating system is the same as for any other boot method.
  - Locate the install image
  - Map it to the virtual CD
  - Boot and install the operating system.
- Access to the storage array console/management interface is optional.

# iSCSI Boot – Booting

## Attach Media



- In the UCS KVM, map the ISO file/CD-ROM to the virtual media.
  - Click on the **Virtual Media** tab in the **KVM Console**
  - Leave the Virtual Media Session up for the duration of the install
  - (Optional) Expand the **Details** button to see the read activity as the media is loaded.
- Just below the menus in the KVM window, click **Boot Server** or **Reset** if the server was already powered up.

# iSCSI Boot – Booting


## Target Login

```
Cisco VIC iSCSI, Boot Driver Version 2.0(3b)
(C) 2010 Cisco Systems, Inc.
0025b501004e iSCSI NETAPP :000
Option ROM installed successfully
```

- The Cisco VIC will log into the target filer.
- The first number after the copyright is the MAC address of the iSCSI vNIC.
- The LUN that was found will be shown.
  - Typically, the boot LUN is zero.
- The login message will show twice, once for each path.

# iSCSI Boot – Booting

## Select Target Disk



- After the installer starts, select the install location.
- Select the NetApp LUN that is seen via iSCSI.
  - The FUJITSU or SEAGATE drives are the local drives on the blade.
  - The LSILOGIC is a logical drive created by the RAID policy on the blade.

# iSCSI Boot – Booting

## LUN Validation

```

                                Disk Details
* Con Model/Vendor:  NETAPP  LUN
----- Full Disk Name:  naa.60a98000503365634b346c704a654f45
Stora LUN ID:          0          acity
----- Target ID:    0          -----
Local Capacity:      10.00 GiB
(n Path:             /vmfs/devices/disks/naa.60a98000503365634b
Remot                346c704a654f45
FU ESX(i) Found:    No          9 GiB
LS Datastores:      (none)     9 GiB
NE                  0 GiB

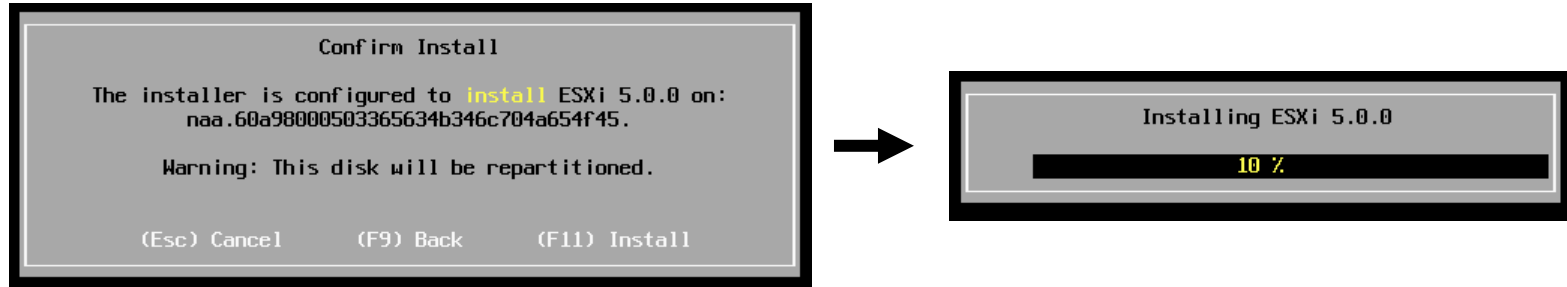
                                (Enter) OK

(Esc) Cancel  (F1) Details  (F5) Refresh  (Enter) Continue
```

- (Optional) Select **<F1>** to view the details of the LUN and whether there is any existing data on the LUN.
  - Press **<Enter>** again to exit and then **<Enter>** again to continue.

# iSCSI Boot – Booting


## Confirm Install



- After selecting a keyboard and setting the system password, confirm the installation onto the iSCSI LUN
- Monitor the installation.

# iSCSI Boot – Booting

## Dedicated iSCSI NICs – Complete



- The installer will *virtually eject* the virtual media.
  - In the KVM Virtual Media the ISO will show as unmapped automatically.

# iSCSI Boot – Booting

## Dedicated iSCSI NICs – Reboot into Operating System

```
Rebooting Server
The server will shut down and reboot.
The process will take a short time to complete.
```

- The server will reboot into the newly installed operating system.



# iSCSI Boot – Booting

## Booted Host



- The host will try to connect via DHCP.

# iSCSI Boot UCS-B with Cisco VIC ESX Configuration



Planning and Deployment

# iSCSI Boot – ESX

## Multipathing

- To ensure path redundancy, add the second iSCSI vNIC into a redundancy group with the primary iSCSI vNIC.

The image shows two screenshots from the VMware vSphere interface. The left screenshot is the 'iScsiBootPG Properties' dialog box, specifically the 'NIC Teaming' tab. It shows a table of network adapters configured for redundancy. A blue box highlights the 'Active Adapters' section, and a blue arrow points from this section to the right screenshot.

Name	Speed	Networks
<b>Active Adapters</b>		
vmnic2	10000 Full	10.86.79.1-10.86.79.127
<b>Standby Adapters</b>		
vmnic3	10000 Full	10.86.79.1-10.86.79.127
<b>Unused Adapters</b>		

The right screenshot shows the 'Configuration' tab for a vSphere Standard Switch. It displays two switch configurations: 'vSwitch0' and 'iScsiBootvSwitch'. The 'iScsiBootvSwitch' configuration shows a 'VMkernel Port' named 'iScsiBootPG' connected to physical adapters 'vmnic2' and 'vmnic3', both at 10000 Full speed. A blue arrow points from the 'Active Adapters' table in the left screenshot to the 'iScsiBootPG' port in this screenshot.

# iSCSI Boot UCS-B with Cisco VIC

## Troubleshooting



Planning and Deployment

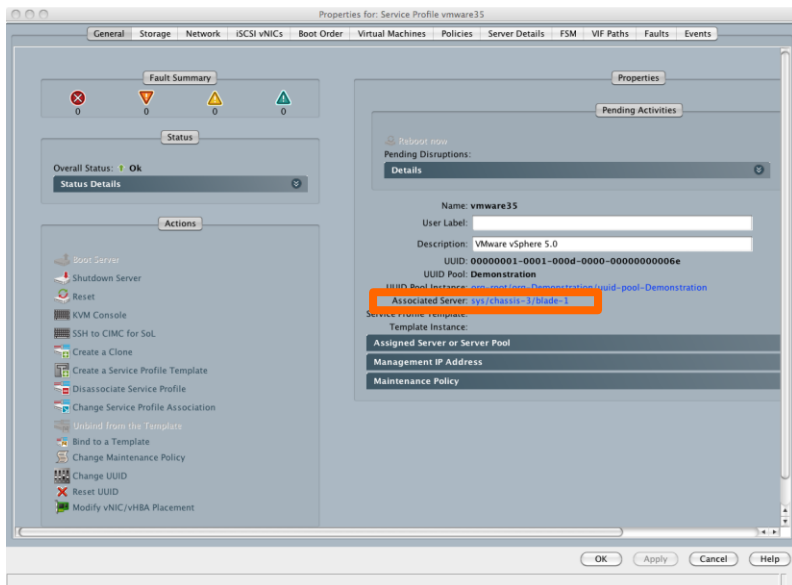
# iSCSI Boot – Troubleshooting

## Basics – VIC Tools

- The Cisco VIC has two built in troubleshooting tools that can be used if the VIC does not find a LUN.
- The tools are available in the UCS Manager CLI.
  - **iscsi\_ping**
    - The ping command will connect to the filer on the configured iSCSI port.
    - The command is only available when the system is booting or has failed to find a boot target.
  - **iscsi\_get\_config**
    - Displays the configuration information programmed on the VNICs as well as the last status.

# iSCSI Boot – Troubleshooting

## Basics – VIC Tools



```
UCS-Prod-A# show service-profile assoc | grep  
-A 1 vmware35  
Demonstration/vmware35  
Associated 3/1
```

- Determine the physical blade the service profile is running on using the GUI or SSH to the CLI.

# iSCSI Boot – Troubleshooting

## Basics – VIC Tools

- Connect to the adapter.
- Use **connect** to get into the VIC's manager.
- Use **attach-mcp** to connect to the master control program.
- From here the iSCSI commands can be used.

```
UCS-Prod-A# connect adapter 3/1/1
adapter 3/1/1 # connect
No entry for terminal type "xterm-color";
using dumb terminal settings.
adapter 3/1/1 (top):1# attach-mcp
No entry for terminal type "xterm-color";
using dumb terminal settings.
adapter 3/1/1 (mcp):1# iscsi<tab>
iscsi_get_config iscsi_ping
```

- ⓘ The **<?>** command does not provide help, use the **<Tab>** key.
- ⓘ The session to the card is disconnected when the blade is rebooted.

# iSCSI Boot – Troubleshooting

## Basics – VIC Tools – iSCSI Ping

```
Adapter 3/1/1 (mcp):7# iscsi_ping
id  name          tgt          address      port  tcp ping status
---  -
  7  vnic_3         0           10.86.79.6   3260  Success (9.875ms)
```

- In the MCP, issue **iscsi\_ping**
- The adapter uses the settings from the service profile to create a TCP connection to the iSCSI port.
  - This verifies/measures connectivity and establishes that iSCSI is enabled on the storage target.
- The target should be the filer and the port being used for iSCSI.
  - If they are not, the target is not defined properly in the service profile.



# iSCSI Boot – Troubleshooting

## Basics – VIC Tools – iSCSI Ping

```
Adapter 3/1/1 (mcp):1# iscsi_ping  
No iSCSI boot vNICs ready.
```

This command can only be run when the host is in the BIOS menu.

- iSCSI Ping is only available during initial boot.

# iSCSI Boot – Troubleshooting

## Basics – VIC Tools – iSCSI Get Config

- In MCP issue **iscsi\_get\_config**.
- Note the configuration and the Status in the initiator section.
  - The initiator is the UCS blade.
- The target is the storage array.
  - Note the Ping State.
  - The VLAN will show 0 despite what is configured.

```
adapter 3/1/1 (mcp):1# iscsi_get_config
```

```
vnic iSCSI Configuration:
```

```
-----
```

```
vnic_id: 7
```

```
link state: Up
```

```
Initiator Cfg:
```

```
initiator_state: ISCSI_INITIATOR_READY
```

```
initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
```

```
vlan: 0
```

```
dhcp status: false
```

```
IQN: iqn.1984-12.com.cisco:vmware-a:35
```

```
IP Addr: 10.86.79.58
```

```
Subnet Mask: 255.255.255.128
```

```
Gateway: 10.86.79.1
```

```
Target Cfg:
```

```
Target Idx: 0
```

```
State: ISCSI_TARGET_READY
```

```
Prev State: ISCSI_TARGET_DISABLED
```

```
Target Error: ISCSI_TARGET_NO_ERROR
```

```
IQN: iqn.1992-08.com.netapp:sn.135077524
```

```
IP Addr: 10.86.79.6
```

```
Port: 3260
```

```
Boot Lun: 0
```

```
Ping Stats: Success (9.338ms)
```

# iSCSI Boot – Troubleshooting

## Failure Scenarios – Bad IQN

```
Cisco VIC iSCSI, Boot Driver Version 2.0(3b)
(C) 2010 Cisco Systems, Inc.
Initialize error 1
```

### UCS KVM

```
ntap01> Wed Aug 15 16:37:13 EST
[ntap01:
iscsi.textkey.invalidValue:warning]:
ISCSI: Invalid value for iSCSI text
key, 'InitiatorName'
```

### Network Appliance Console

- Error1 indicates problems finding the LUN.
  - The NetApp console is the best source of information in this case.
- The source IQN is not a valid format.

```
adapater 3/1/1 (mcp):1# iscsi_get_config
```

```
vnic iSCSI Configuration:
-----
```

```
vnic_id: 5
    link_state: Up

    Initiator Cfg:
        initiator_state: ISCSI_INITIATOR_READY
        initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
        vlan: 0
        dhcp status: false
        IQN: esx-demo-vmware-a:35
        IP Addr: 10.86.79.58
        Subnet Mask: 255.255.255.128
        Gateway: 10.86.79.1
```

```
    Target Cfg:
        Target Idx: 0
        State: INVALID
        Prev State: ISCSI_TARGET_GET_SESSION_INFO
        Target Error: ISCSI_TARGET_LOGIN_ERROR
        IQN: iqn.1992-08.com.netapp:sn.135077524
        IP Addr: 10.86.79.6
        Port: 3260
        Boot Lun: 0
        Ping Stats: Success (9.966ms)
```

### UCS VIC MCP Console

# iSCSI Boot – Troubleshooting

## Failure Scenarios – No Initiator Map or LUN Map

```
Cisco VIC iSCSI, Boot Driver Version 2.0(3b)
(C) 2010 Cisco Systems, Inc.
Initialize error 1
```

UCS KVM

```
ntap01(takeover)> Wed Aug 15 17:26:53 EST
[ntap01 (takeover): iscsi.notice:notice]:
ISCSI: New session from initiator
iqn.1984-12.com.cisco:vmware-a:35 at IP
addr 10.86.79.58
```

Network Appliance Console

- The service console shows a valid connection but it does not have an initiator mapping or the mapping goes to the wrong LUN number.
  - Use **lun map -v** on the filer to validate the map.

```
adapter 3/1/1 (mcp):1# iscsi_get_config
```

```
vnuc iSCSI Configuration:
-----
```

```
vnuc_id: 5
      link_state: Up

Initiator Cfg:
  initiator_state: ISCSI_INITIATOR_READY
  initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
      vlan: 0
  dhcp status: false
      IQN: iqn.1984-12.com.cisco:vmware-a:35
  IP Addr: 10.86.79.58
  Subnet Mask: 255.255.255.128
  Gateway: 10.86.79.1

Target Cfg:
  Target Idx: 0
      State: INVALID
  Prev State: ISCSI_TARGET_GET_LUN_INFO
  Target Error: ISCSI_TARGET_GET_HBT_ERROR
      IQN: iqn.1992-08.com.netapp:sn.135077524
  IP Addr: 10.86.79.6
  Port: 3260
  Boot Lun: 0
  Ping Stats: Success (9.388ms)
```

UCS VIC MCP Console

# iSCSI Boot – Troubleshooting

## Failure Scenarios – No Initiator Map or LUN Map – LUN Validation

```
ntap01> lun show -v all
  /vol/FV_Demo_LUNs/vmware35      10g (10737418240) (r/w, online, mapped)
    Comment: "UCS-Prod / Demonstration Org"
    Serial#: P3ecK4/nvXC3
    Share: none
    Space Reservation: enabled
    Multiprotocol Type: linux
    Maps: vmware35=0
```

- Verify the LUN size is big enough for the operating system
- Verify the path is online and read/write
- Verify the protocol for the boot LUN is Linux
- Verify the LUN Mapping

# iSCSI Boot – Troubleshooting

## Failure Scenarios – No Operating System

```
Cisco VIC iSCSI, Boot Driver Version 2.0(3b)
(C) 2010 Cisco Systems, Inc.
 0025b501004e iSCSI NETAPP :000
Option ROM installed successfully
```

```
Reboot and Select proper Boot device
or Insert Boot Media in selected Boot device and press a key
```

### UCS KVM

```
ntap01> Wed Aug 15 17:44:04 EST [ntap01:
iscsi.notice:notice]: ISCSI: New session from
initiator iqn.1984-12.com.cisco:vmware-a:35 at IP
addr 10.86.79.58
Wed Aug 15 17:44:19 EST [ntap01:
iscsi.notice:notice]: ISCSI: New session from
initiator iqn.1984-12.com.cisco:vmware-b:35 at IP
addr 10.86.79.59
```

### Network Appliance Console

- If the option rom shows success, the LUN has been found.
  - Validate the operating system is not corrupt/is installed.

```
adapter 3/1/1 (mcp):1# iscsi_get_config
vnic_id: 5
    link_state: Up

Initiator Cfg:
    initiator_state: ISCSI_INITIATOR_READY
    initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
    vlan: 0
    dhcp status: false
    IQN: iqn.1984-12.com.cisco:vmware-a:35
    IP Addr: 10.86.79.58
    Subnet Mask: 255.255.255.128
    Gateway: 10.86.79.1

Target Cfg:
    Target Idx: 0
    State: ISCSI_TARGET_READY
    Prev State: ISCSI_TARGET_DISABLED
    Target Error: ISCSI_TARGET_NO_ERROR
    IQN: iqn.1992-08.com.netapp:sn.135077524
    IP Addr: 10.86.79.6
    Port: 3260
    Boot Lun: 0
    Ping Stats: Success (9.493ms)
```

```
Session Info:
    session_id: 0
    host_number: 0
    bus_number: 0
    target_id: 0
```

# iSCSI Boot – Troubleshooting

## Failure Scenarios – Jumbo Frames

```
ntap01> iscsi interface show
Interface e0a enabled
Interface e0b enabled
Interface e1 enabled

ntap01> ifconfig -a
e0a: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCSUM> mtu 1500
    ether 00:a0:98:24:0c:f4 (auto-unknown-cfg_down) flowcontrol full
e0b: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCSUM> mtu 1500
    ether 00:a0:98:24:0c:f5 (auto-unknown-cfg_down) flowcontrol full
e1: flags=0x3d48867<UP,BROADCAST,RUNNING,MULTICAST,TCPCSUM> mtu 9000
    inet 10.86.79.6 netmask 0xffffffff80 broadcast 10.86.79.127
    ether 00:07:43:06:25:e3 (auto-10g_sr-fd-up) flowcontrol full
lo: flags=0x1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCSUM> mtu 8160
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1
    ether 00:00:00:00:00:00 (VIA Provider)
```

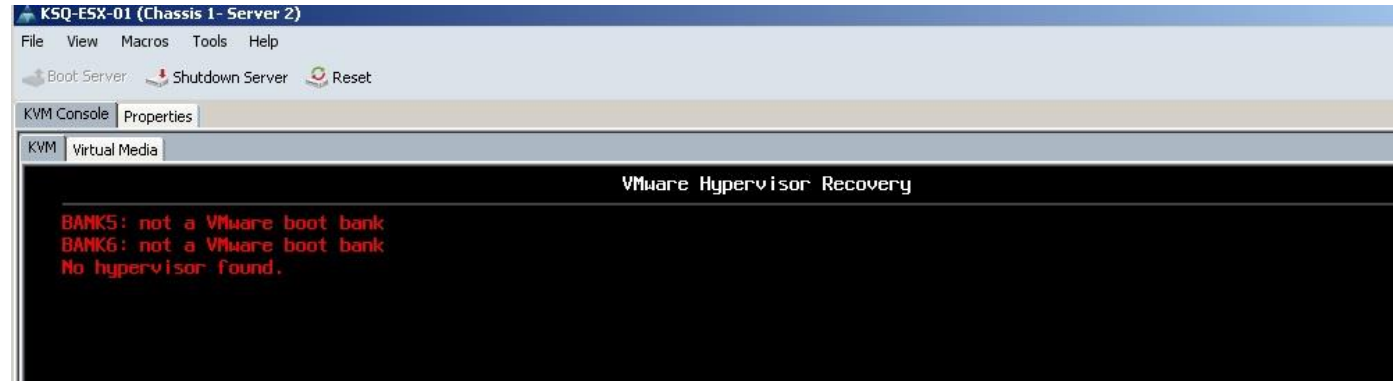
- If the connection succeeds but the OS fails to load, verify the Jumbo Frames are properly enabled.
  - On the filer, verify the MTU is set to 9000 for the interface iSCSI is using.

# iSCSI Boot – Troubleshooting

Failure Scenarios – Jumbo Frames – ESXi complaints

```
Kernel: b.z bootUUID=7b477d5f2c9e7e98b6144ac322fe4b40
Loading b.z
PANIC: Error while reading file: -1, b.z
boot:
```

Esxi 4.1



The screenshot shows the VMware vSphere interface for a host named 'KSQ-ESX-01 (Chassis 1- Server 2)'. The 'KVM Console' tab is active, displaying a 'VMware Hypervisor Recovery' error. The error text is as follows:

```
BANK5: not a VMware boot bank
BANK6: not a VMware boot bank
No hypervisor found.
```

Esxi 5.0



# iSCSI Boot – Troubleshooting Logins

## IQN Validation – Initiator

```
ntap01> igroup show
```

```
vmware35 (iSCSI) (ostype: linux):
```

```
iqn.1984-12.com.cisco:vmware-a:35 (not logged in)
```



```
adapter 3/1/1 (mcp):1# iscsi_get_config
```

```
...
```

```
Initiator Cfg:
```

```
initiator_state: ISCSI_INITIATOR_READY
```

```
initiator_error_code: ISCSI_BOOT_NIC_NO_ERROR
```

```
vlan: 0
```

```
dhcp status: false
```

```
IQN: iqn.1984-12.com.cisco:vmware-a:35
```

```
IP Addr: 10.86.79.58
```

```
Subnet Mask: 255.255.255.128
```

```
Gateway: 10.86.79.1
```

- Verify the IQN must match the configuration on the VIC. It must match exactly.

# iSCSI Boot – Troubleshooting Logins


## iQN Validation – Target

```
ntap01> iscsi nodename
```


```
iSCSI target nodename: iqn.1992-08.com.netapp:sn.135077524
```

```
ntap01> lun show -m
```

LUN path	Mapped to	LUN ID	Protocol
/vol/FV_Demo_LUNs/vmware31	vmware31	0	iSCSI
...			



```
adapter 3/1/1 (mcp):1# iscsi_get_config
Target Cfg:
Target Idx: 0
State: ISCSI_TARGET_READY
Prev State: ISCSI_TARGET_DISABLED
Target Error: ISCSI_TARGET_NO_ERROR
IQN: iqn.1992-08.com.netapp:sn.135077524
IP Addr: 10.86.79.6
Port: 3260
Boot Lun: 0
```



- Verify the configured iSCSI node name matches the filer and the target  
Boot LUN matches the LUN map.

# iSCSI Boot – Troubleshooting

## Login

```
Wed Aug 15 18:58:42 EST [ntap01: iscsi.notice:notice]: ISCSI: New session from initiator  
iqn.1984-12.com.cisco:vmware-a:35 at IP addr 10.86.79.58  
Wed Aug 15 18:58:53 EST [ntap01: iscsi.notice:notice]: ISCSI: New session from initiator  
iqn.1984-12.com.cisco:vmware-b:35 at IP addr 10.86.79.59
```

- On the Netapp console, the login will be shown in real time.
- Note that both the paths will log in using their respective initiator IP addresses.

# iSCSI Boot – Troubleshooting

## Who's Connected?

```
ntap01> iscsi initiator show
Initiators connected:
  TSIH  TPGroup  Initiator/ISID/IGroup
    4    1002    iqn.1995-05.com.broadcom.vmware31.iscsiboot / 00:02:3d:00:00:02 /
vmware31
    35   1002    iqn.1984-12.com.cisco:vmware-a:35 / 00:02:3d:00:00:02 / vmware35
```

- The **iscsi initiator show** command provides a list of active connections.
  - Note how the B paths are not actively in use.

# iSCSI Boot – Troubleshooting

## Connection View

- The active sessions show the settings and if there is active traffic, which commands are being executed.

```
ntap01> iscsi session show -v
Session 35
  Initiator Information
    Initiator Name: iqn.1984-12.com.cisco:vmware-a:35
    ISID: 00:02:3d:00:00:02

  Session Parameters
    SessionType=Normal
    TargetPortalGroupTag=1002
    MaxConnections=1
    ErrorRecoveryLevel=0
    AuthMethod=None
    HeaderDigest=None
    DataDigest=None
    ImmediateData=Yes
    InitialR2T=No
    FirstBurstLength=65536
    MaxBurstLength=65536
    Initiator MaxRecvDataSegmentLength=65536
    Target MaxRecvDataSegmentLength=65536
    DefaultTime2Wait=2
    DefaultTime2Retain=0
    MaxOutstandingR2T=1
    DataPDUInOrder=Yes
    DataSequenceInOrder=Yes
    Command Window Size: 64

  Connection Information...
```

# iSCSI Boot – Troubleshooting

## ESX Configuration Changes Not Persistent

- Configure ESX console, host, DNS, VLAN settings.
- Perform a warm reboot of the host.
- The host comes up with no settings and using the iSCSI initiator IP as the service console.
- There is a known issue in some flavors of VMware ESXi 5.0 with iSCSI and FCoE boot where the configuration is not persistent.
- Resolution: Use ESX 5.0 U2 or 5.1.
  - [http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2031315](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2031315)

### UCS Blade loses configuration after reboot in VMware ESXi 5.0.x

#### Symptoms

- UCS Blades booting from FCoE/iSCSI SAN revert to a previous state after reboot.
- The ESXi 5.0 Bootbank points to /tmp and/or no altbootbank exists.
- You are unable to install VMware Tools in virtual machines.
- You are unable to patch/update ESXi 5.0.
- A task fails with:

BUILT FOR  
THE HUMAN  
NETWORK

