Navigating DHCP Processes in SD-Access Network

Cisco Community

Carlos Nicoletti, Technical Consulting Engineer – SDA/Catalyst Center
Diego Cabañas, Team Captain – SDA/Catalyst Center

Wednesday, November 13th 2024

Our Experts

Carlos Nicoletti

Technical Consulting Engineer

With over 7 years of experience at Cisco. His career began with the Internet of Things (IoT) team and has evolved to his current role in the Cisco Catalyst Center team, specifically focusing on SD Access. His expertise includes topics such as DHCP in Fabric, LISP, PubSub, Plug and Play, LAN Automation, and Cisco Trustsec.

He has also been recognized for his participation in events like Cisco Live and Cisco Connect, where he has shared his expertise and knowledge. Carlos graduated from the Universidad Nacional Autonoma de Mexico (UNAM) and is an ex-professor at this university.

He also holds certifications in CCNA, CCNP Enterprise, ScrumMaster, and ITIL v4.

© 2024 Cisco and/or its affiliates. All rights reserved.

Our Experts

Diego Cabañas

Team Captain

With over a decade of experience in networking. He studied Electronic Technologies Engineering at Tec de Monterrey and joined Cisco in 2014.

Diego is a father of three who enjoys reading and staying updated with the latest innovations. Specializing in Campus Automation for Enterprise Networks,

Diego focuses on switching architectures, SDA, Catalyst Center, and automation technologies. He currently solves complex networking challenges and designs solutions that enhance the efficiency, scalability, and security

© 2024 Cisco and/or its affiliates. All rights reserved.

Join at

slido.com

#2835 026

Passcode: **xnygs2**

Agenda

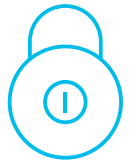

1. Fundamentals of DHCP in SD-Access
2. Limitations of DHCP in SD-Access networks
3. DHCP process within SD-Access
4. Troubleshooting common issues

© 2024 Cisco and/or its affiliates. All rights reserved.

Fundamentals of DHCP in SD-Access

- Fundamentals of DHCP in SD-Access
- Limitations of DHCP in SD-Access networks
- DHCP process within SD-Access
- Troubleshooting common issues

© 2024 Cisco and/or its affiliates. All rights reserved.

What is SD-Access?

- **Built on the principles of intent-based networking, SD-Access helps organizations enable policy-based automation from the edge to the cloud.**
- SD-Access provides network architects with the tools needed to orchestrate key business functions such as user onboarding, secure segmentation, IoT integration, and guest access.
- SD-Access automates user and device policy for any application across both wireless and wired networks through a single fabric network.
- And SD-Access is a transformational change. It allows IT to configure network access in minutes for any user, device, or application, without compromising security.

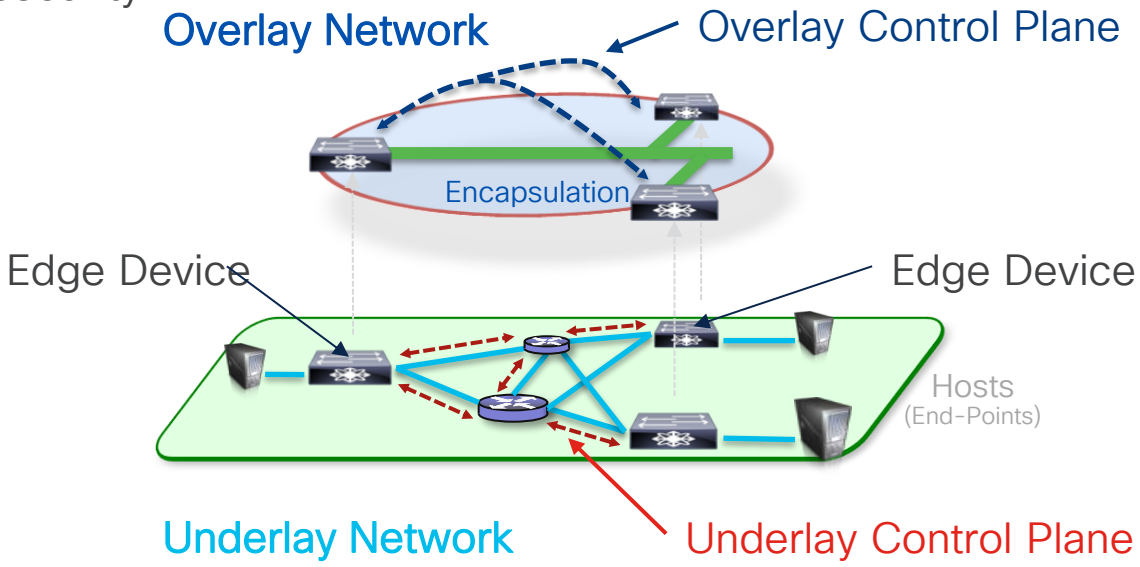

© 2024 Cisco and/or its affiliates. All rights reserved.

Example of an SD-Access Network

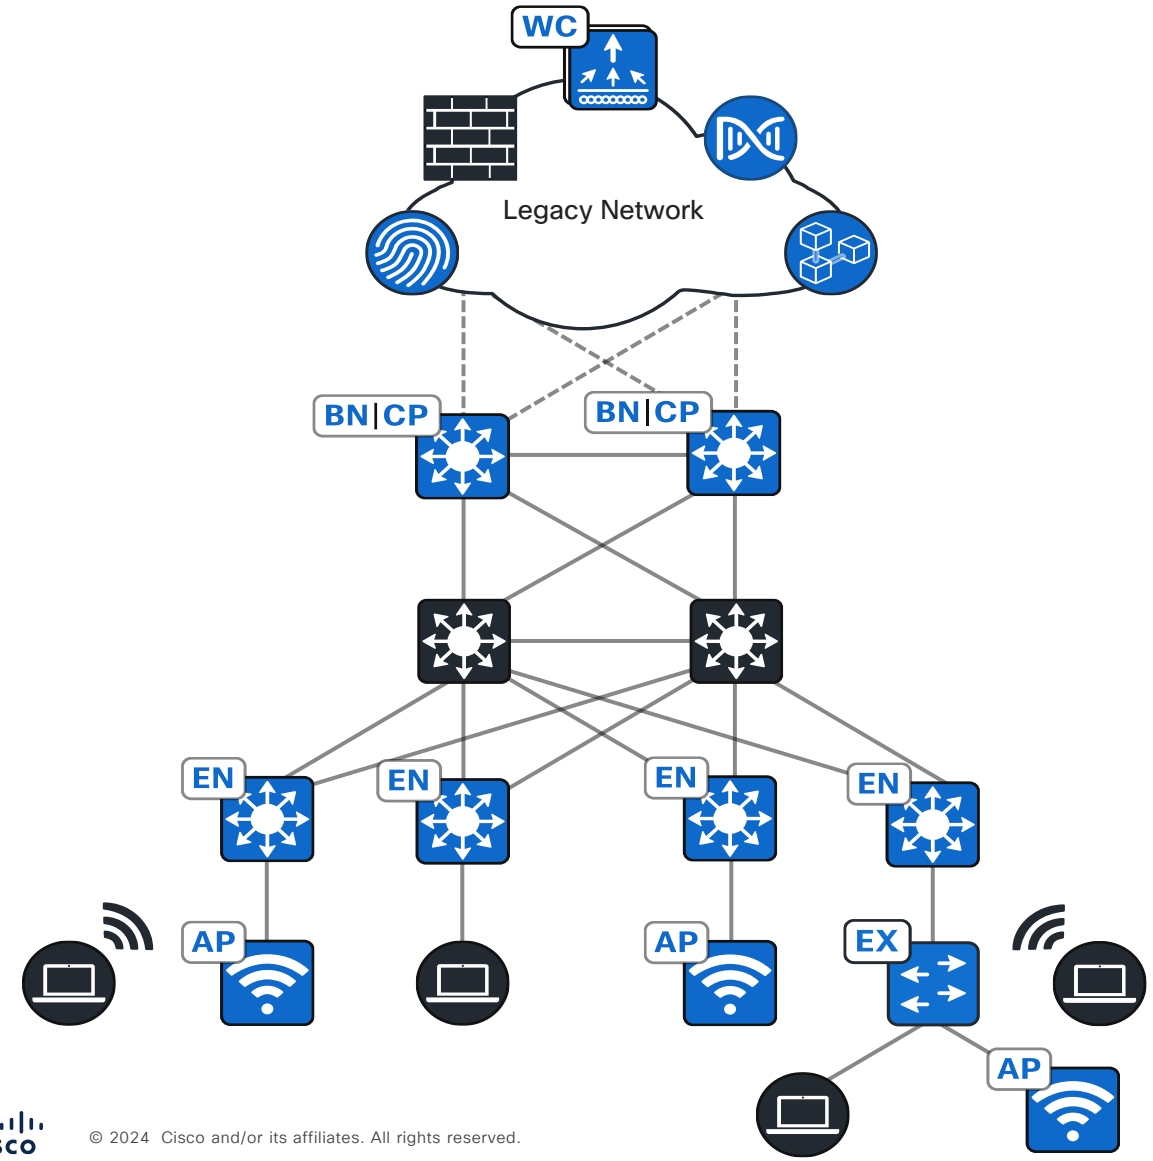

Orchestrator Plane:

Cisco Catalyst Center

Control Plane:

Cisco Locator-ID Separation Protocol (LISP)

Data Plane:

Virtual Extensible LAN (VXLAN)

Security Plane:

Trustsec (SGT, SXP, Inline Tagging)

ISE

© 2024 Cisco and/or its affiliates. All rights reserved.

Anycast Gateway

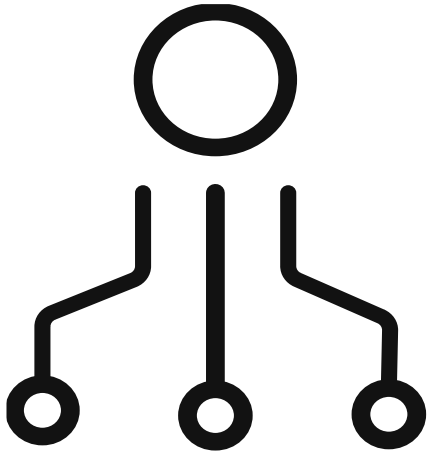

It is a technique that allows a single IP address to be used by multiple devices (or gateways) in a network, and the traffic is automatically routed to the nearest (or most suitable) device that holds that IP address. This technique is based on the principle of anycast, which is a form of addressing in which a single identifier is shared by multiple points of presence.

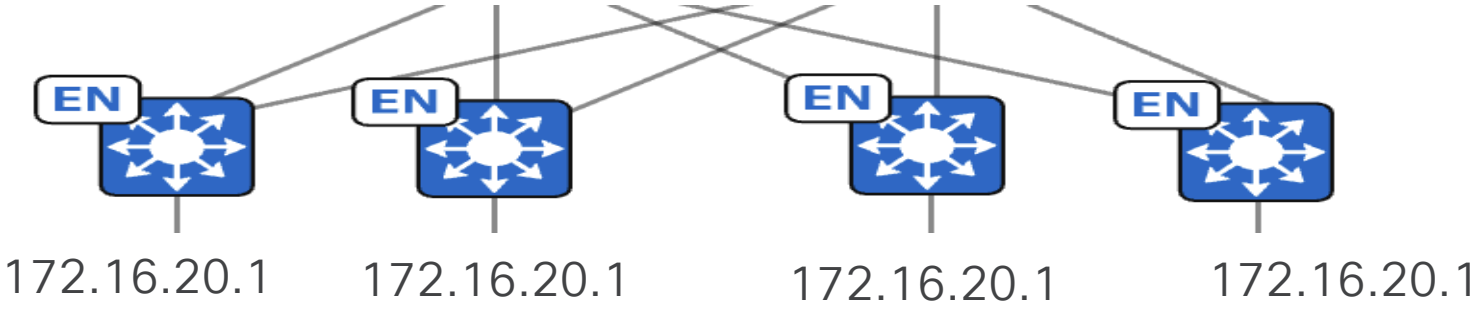

© 2024 Cisco and/or its affiliates. All rights reserved.

What is the DHCP Protocol used for?

The DHCP Protocol (Dynamic Host Configuration Protocol) is used to automatically assign IP addresses and other network parameters to devices (such as computers, phones, printers, etc.) on a network. Its main function is to simplify the configuration of devices connected to a network, avoiding the need for an administrator to manually configure each device with a static IP address.

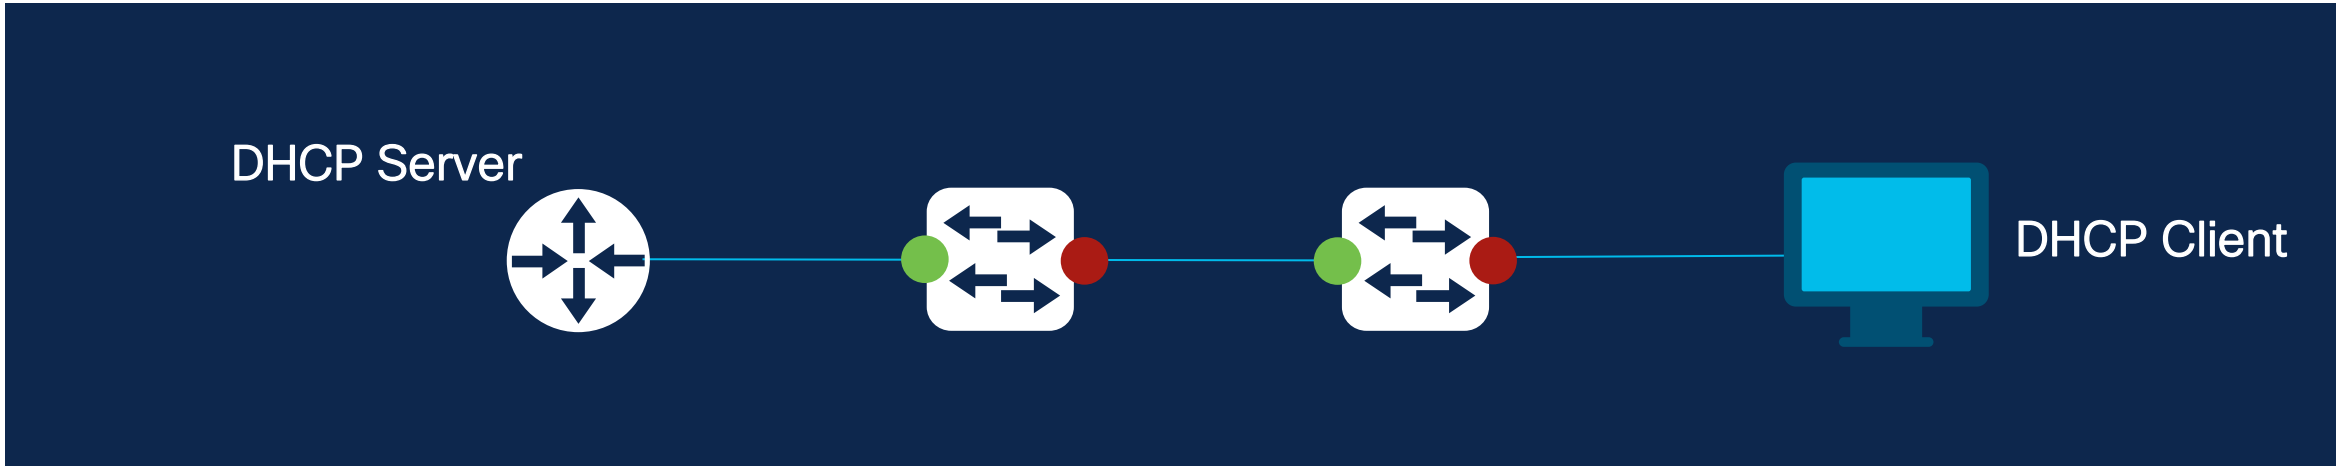

© 2024 Cisco and/or its affiliates. All rights reserved.

Messages Involved in DHCP

Messages Sent by the Client in DHCP:

- DISCOVER
- REQUEST
- RELEASE: Used to tell the sever that client does no longer need an IP address
- DECLINE: Used to decline the IP address offered by a DHCP server

Messages sent by DHCP server:

- OFFER
- ACK
- NAK: used to decline a DHCP request message

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP Process to Obtain an IP Address

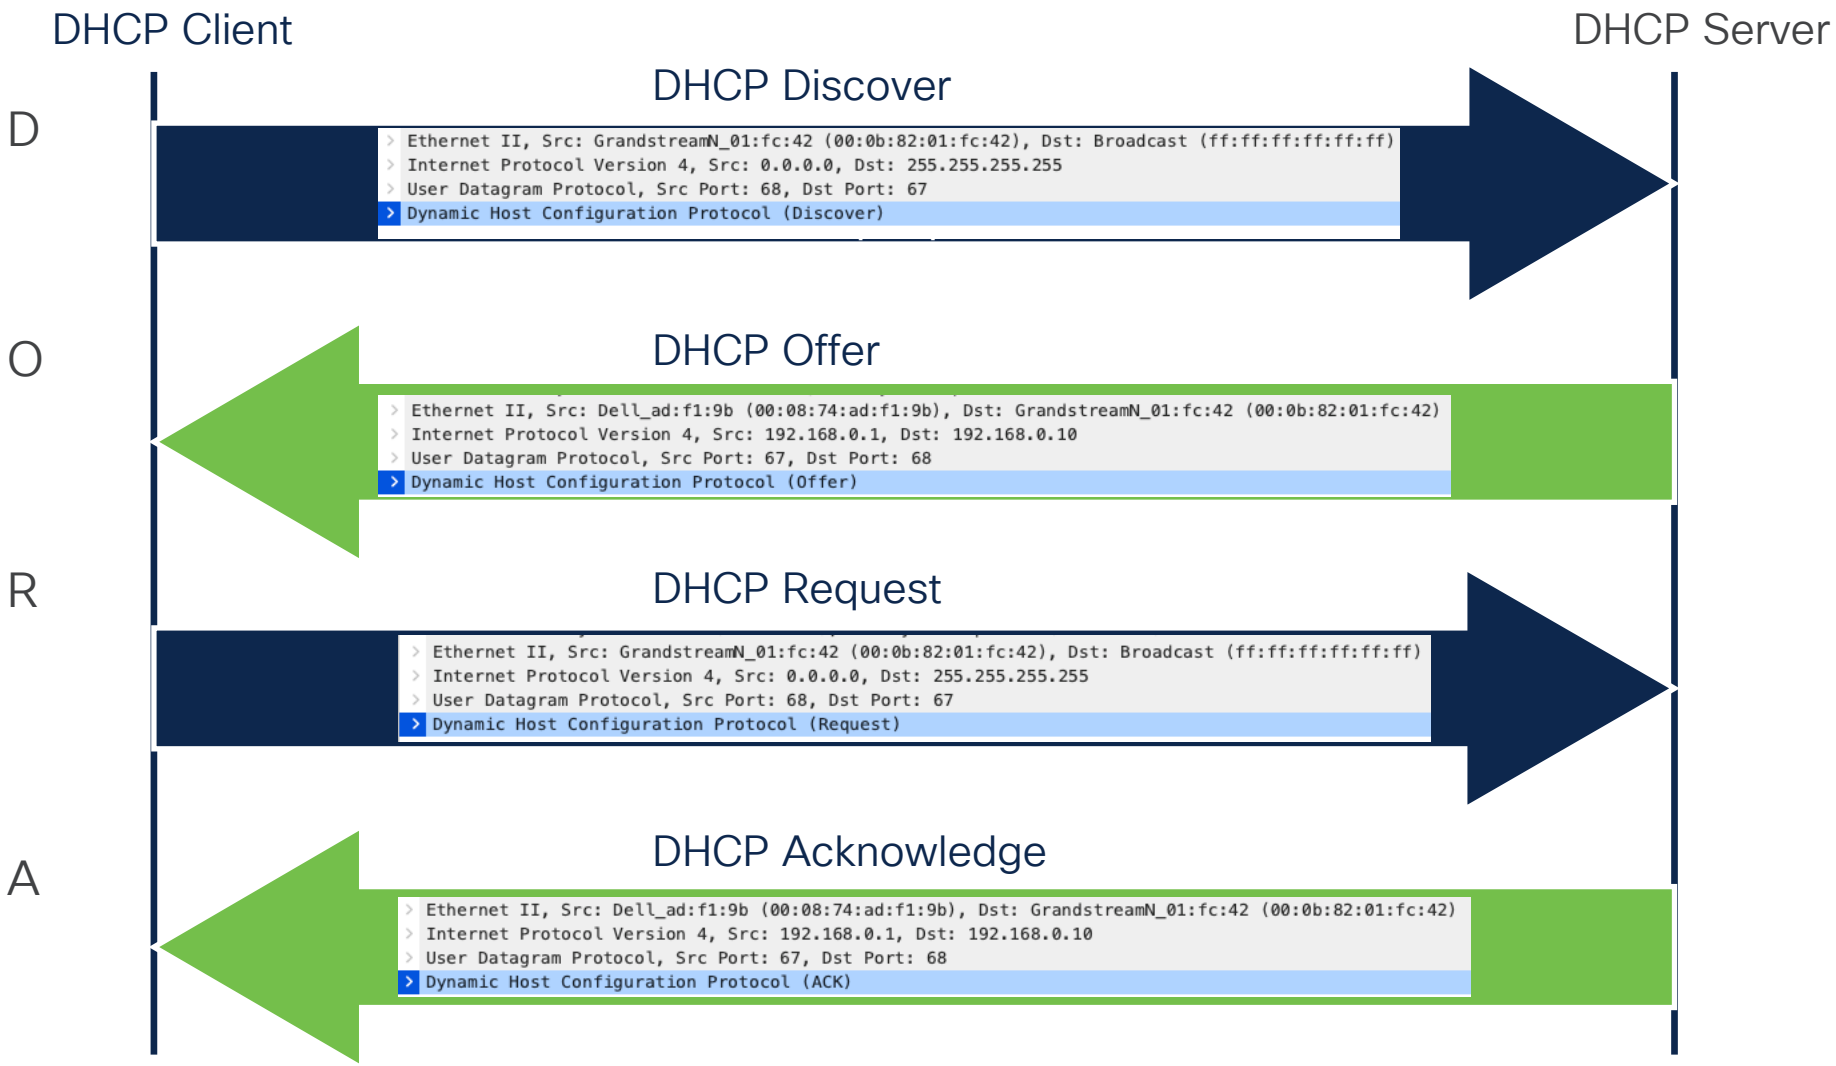

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP Relay

- It is a mechanism used to forward DHCP messages between clients and servers when they are not on the same local network (subnet). In scenarios where a DHCP server is located on a different subnet than the client, a DHCP relay agent is used to pass DHCP messages between the client and server. The **ip helper-address command** enables DHCP Relay Agent functionality in Cisco IOS-XE.

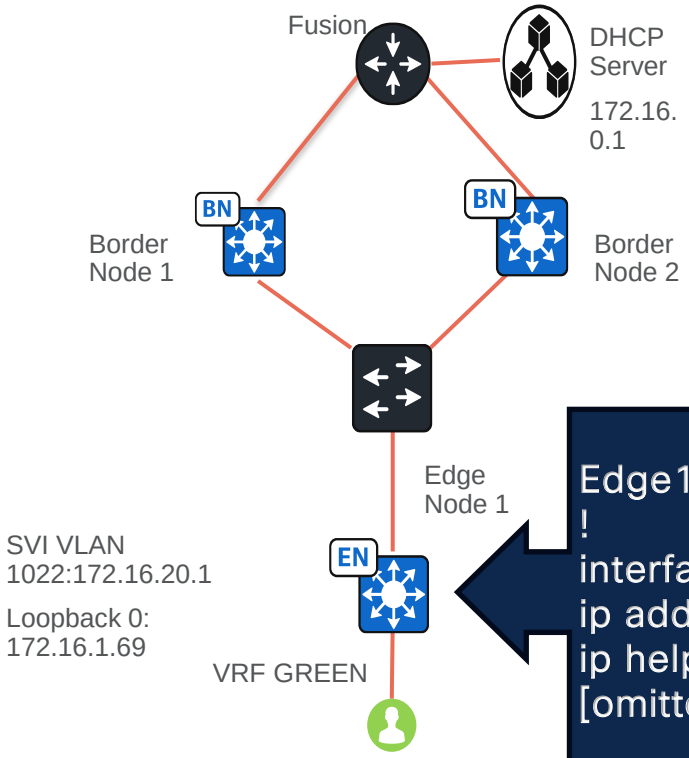


```
Edge1#show run interface vlan 1022
!
interface Vlan1022
ip address 172.16.20.1 255.255.255.0
ip helper-address 172.16.0.1
[omitted]
```

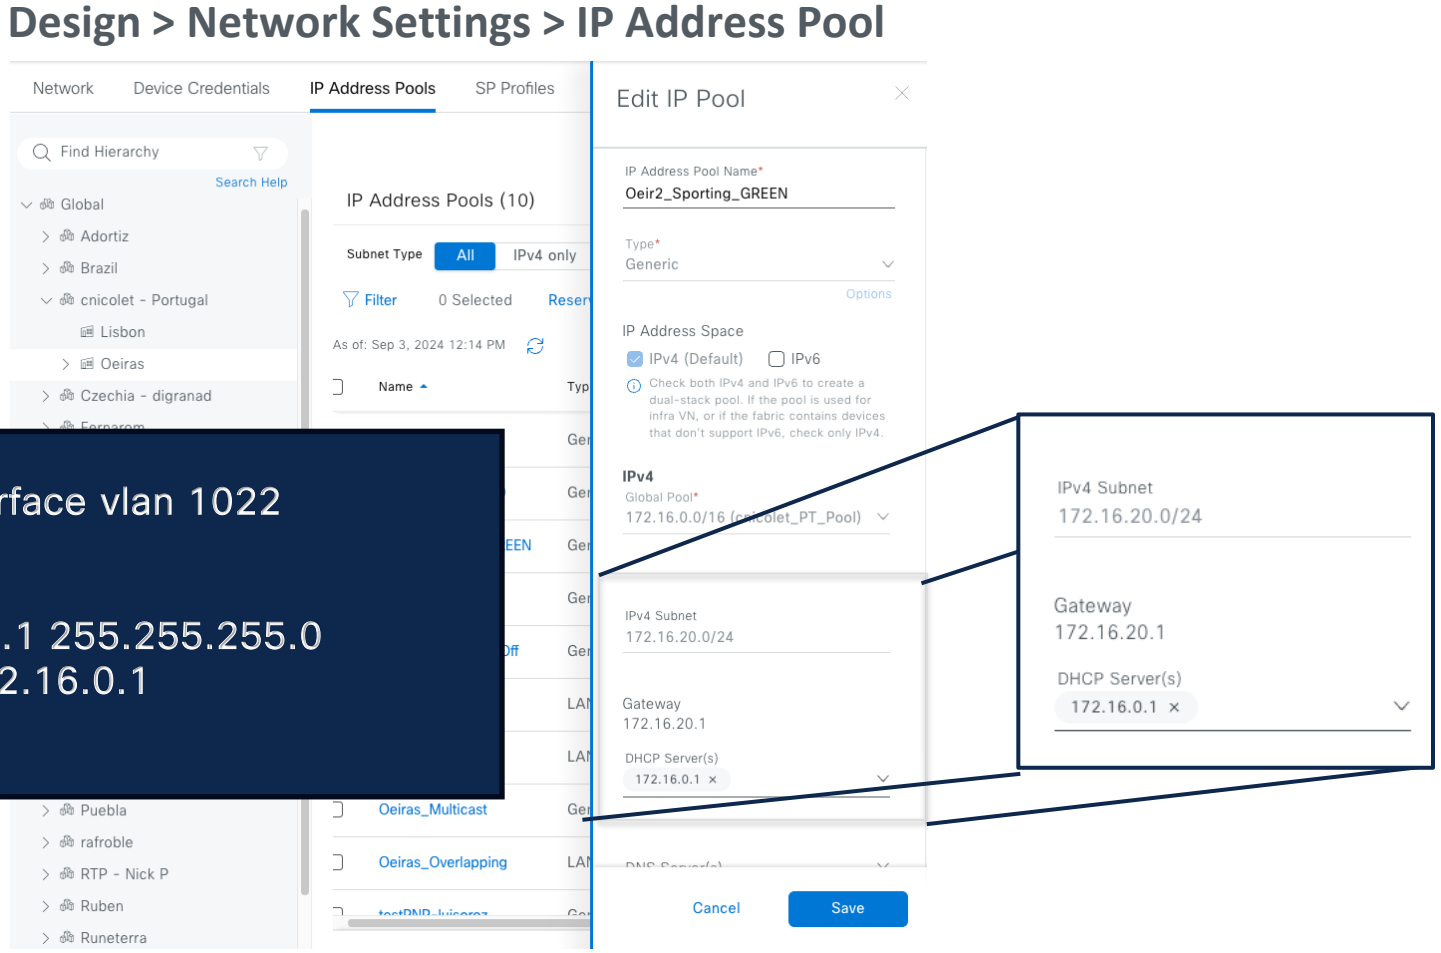

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP option 82 DHCP Relay Agent Information Option

- It is an extension of the DHCP protocol that provides additional information to DHCP servers about the location of the client requesting an IP address. This option is particularly useful in large or complex networks where clients are connected through multiple network devices such as switches and routers.

```
…
Dynamic Host Configuration Protocol (Discover)
    Message type: Boot Request (1)
…
Option: (82) Agent Information Option
        Length: 20
Option 82 Suboption: (1) Agent Circuit ID
            Length: 6
            Agent Circuit ID: 000403fe010a
Option 82 Suboption: (2) Agent Remote ID
            Length: 10
            Agent Remote ID: 030800100501ac100145…
```

© 2024 Cisco and/or its affiliates. All rights reserved.

Who generates DHCP Option 82?

DHCP Snooping is the feature that helps insert Option 82.

```
Circuit ID value analyzed:  000403fe010a
Remote ID value analyzed:  030800100501ac100145

From Circuit ID:
---------
VLAN: 1022                      hex = 0x3fe
Module:  1                      hex = 0x01
Port:  10                       hex = 0x0a


From Remote ID:
---------
Sub-option: 3                   hex = 0x03
Length of option: 8             hex = 0x08
LISP Instance ID: 4101          hex = 0x01005
IP Locator : IPv4               hex = 0x01
Source Locator: 172.16.1.69     hex = 0xac100145 ← Loopback0 IP
```

```
Switch# show ip dhcp snooping

Switch DHCP snooping is enabled

DHCP snooping is configured on following VLANs:

11,101,1021-1022,1025-1026

DHCP snooping is operational on following VLANs:

11,101,1021-1022,1025-1026

 [snipped]

Insertion of option 82 is enabled

   circuit-id default format: vlan-mod-port

   remote-id: 00bf.7798.4980 (MAC)

Option 82 on untrusted port is not allowed

Verification of hwaddr field is enabled

Verification of giaddr field is enabled
```

© 2024 Cisco and/or its affiliates. All rights reserved.

What is DHCP Snooping and How it Works?

DHCP Snooping is a Layer 2 security feature on switches that acts as a "trusted" or "untrusted" filter for DHCP traffic.

Configuration of Trusted and Untrusted Ports:

- **Trusted Ports (trust):** These are the switch ports that are connected to legitimate DHCP servers and other switches or routers that are expected to act as DHCP servers. DHCP snooping will not inspect messages on trusted ports; the switch will simply forward them normally.
- Untrusted Ports (untrust): **These are the ports connected to end devices, such as workstations and computers. These ports are configured as untrusted. These ports should not send DHCP server messages. DHCP snooping will inspect DHCP messages on untrusted ports.**

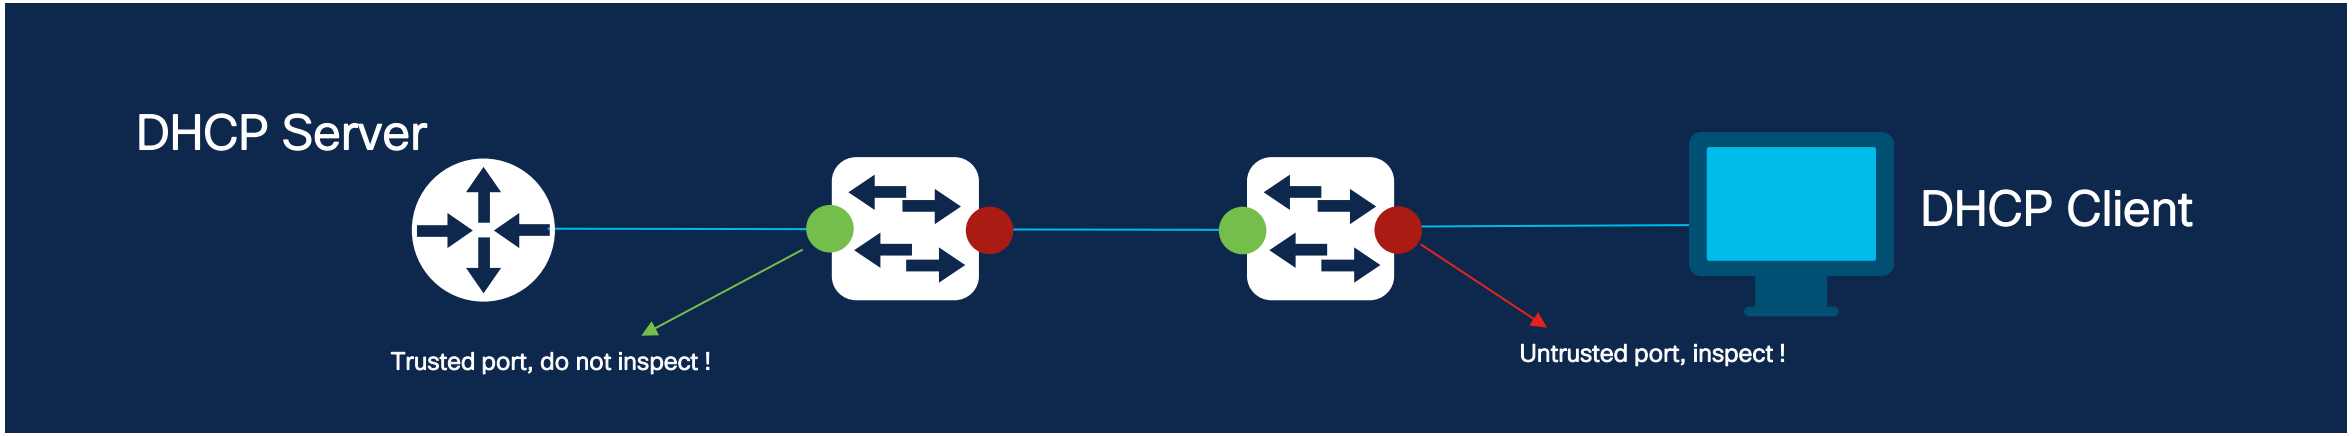

© 2024 Cisco and/or its affiliates. All rights reserved.

What is DHCP Snooping and How It Works?

DHCP Message Filtering:

Allowed Messages: Only DHCP server messages (DHCPOFFER, DHCPACK) are allowed through trusted ports.

Blocked Messages: DHCP request messages (DHCPREQUEST) and other messages from untrusted ports are blocked, except for DHCPDISCOVER and DHCPREQUEST messages from clients.

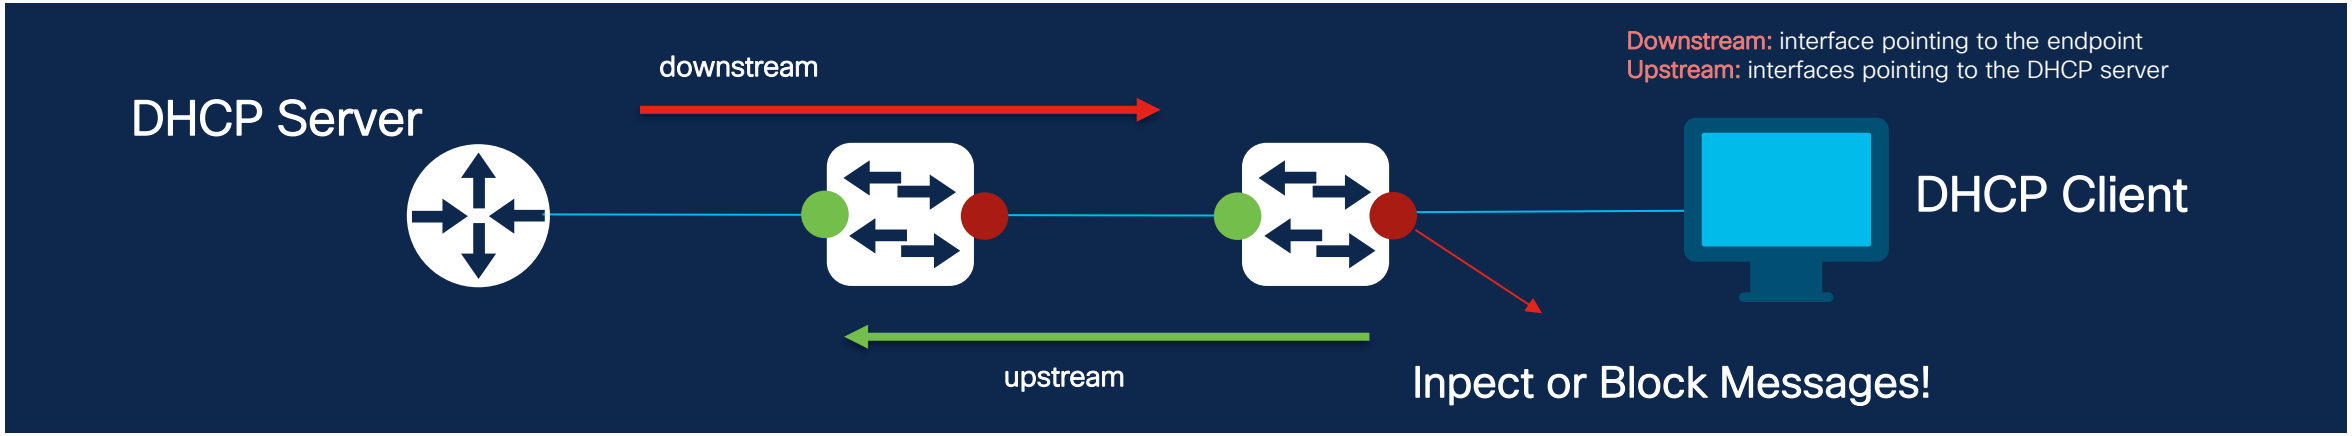

© 2024 Cisco and/or its affiliates. All rights reserved.

Limitations of DHCP in SD-Access networks

© 2024 Cisco and/or its affiliates. All rights reserved.

Limitations of DHCP in SD-Access Networks

- **DHCP servers cannot connect to a Fabric Edge as part of a Fabric IP Pool.**
  - Any DHCP packet received with VXLAN encapsulation destined for anything other than the Anycast gateway IP address will be discarded.
  - There is an **exception** when we have an L2-only VLAN in the Fabric.

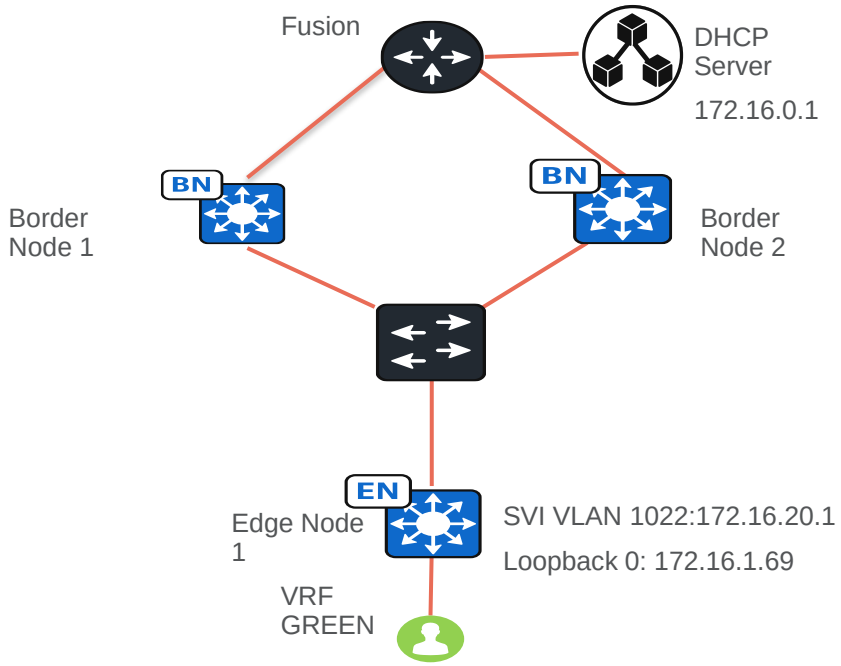

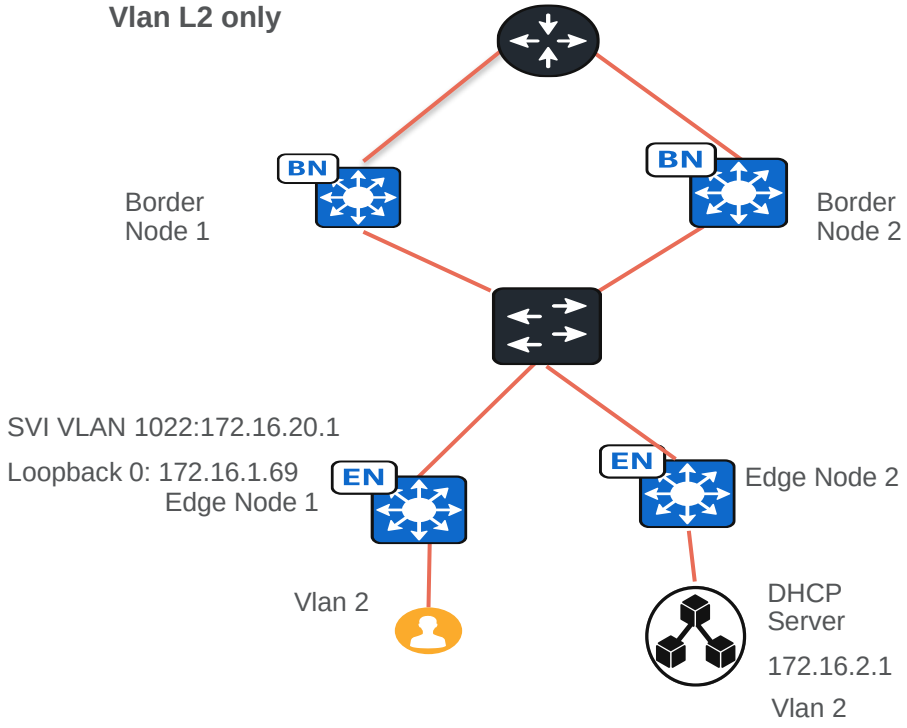

© 2024 Cisco and/or its affiliates. All rights reserved.

Limitations of DHCP in SD-Access Networks

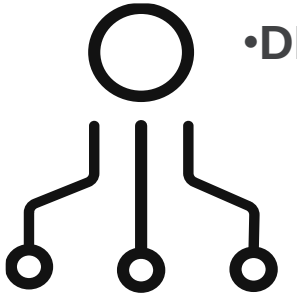

- **DHCP servers must accept and retain the DHCP Option 82.**
  - DHCP Relay Agents will use the source IP of the SVI (Anycast Gateway) of each IP Pool to forward the packet. This requires additional information to identify the correct Edge/Relay Agent.

- **Avoid disabling DHCP snooping as a workaround when validating issues.**
  - DHCP snooping is the component responsible for inserting Option 82 into the DHCP Discover packet; without this Option 82 present, DHCP Offer packets will not be sent to the correct switch.

© 2024 Cisco and/or its affiliates. All rights reserved.

Join at

slido.com

#2835 026

Passcode:

xnygs2

It is possible to have a DHCP server inside the Fabric?

True

0%

Just with L3 Only

0%

False

0%

Just with L2 Only

0%

Limitations of DHCP in SD-Access Networks

- **Bidirectional applications may not work with relays (IP helpers).**
  - The addresses specified in IP helpers can relay a variety of UDP protocols; most of them, except for DHCP, do not include a value like Option 82 to identify the source of the packet.

TAC Tip

The following UDP port combinations are allowed for processing DHCP packets:

- Source Port: 68, Destination Port: 67
- Source Port: 67, Destination Port: 67
- Source Port: 67, Destination Port: 68

Any other port combination will be DROPPED.

© 2024 Cisco and/or its affiliates. All rights reserved.

Join at

slido.com

#2835 026

Passcode:

xnygs2

Which is the necessary option in DHCP to operate correctly in the SD-Access Fabric?

Option 66

0%

Option 82

0%

Option 43

0%

Option 67

0%

DHCP process within SD-Access

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access- Discover

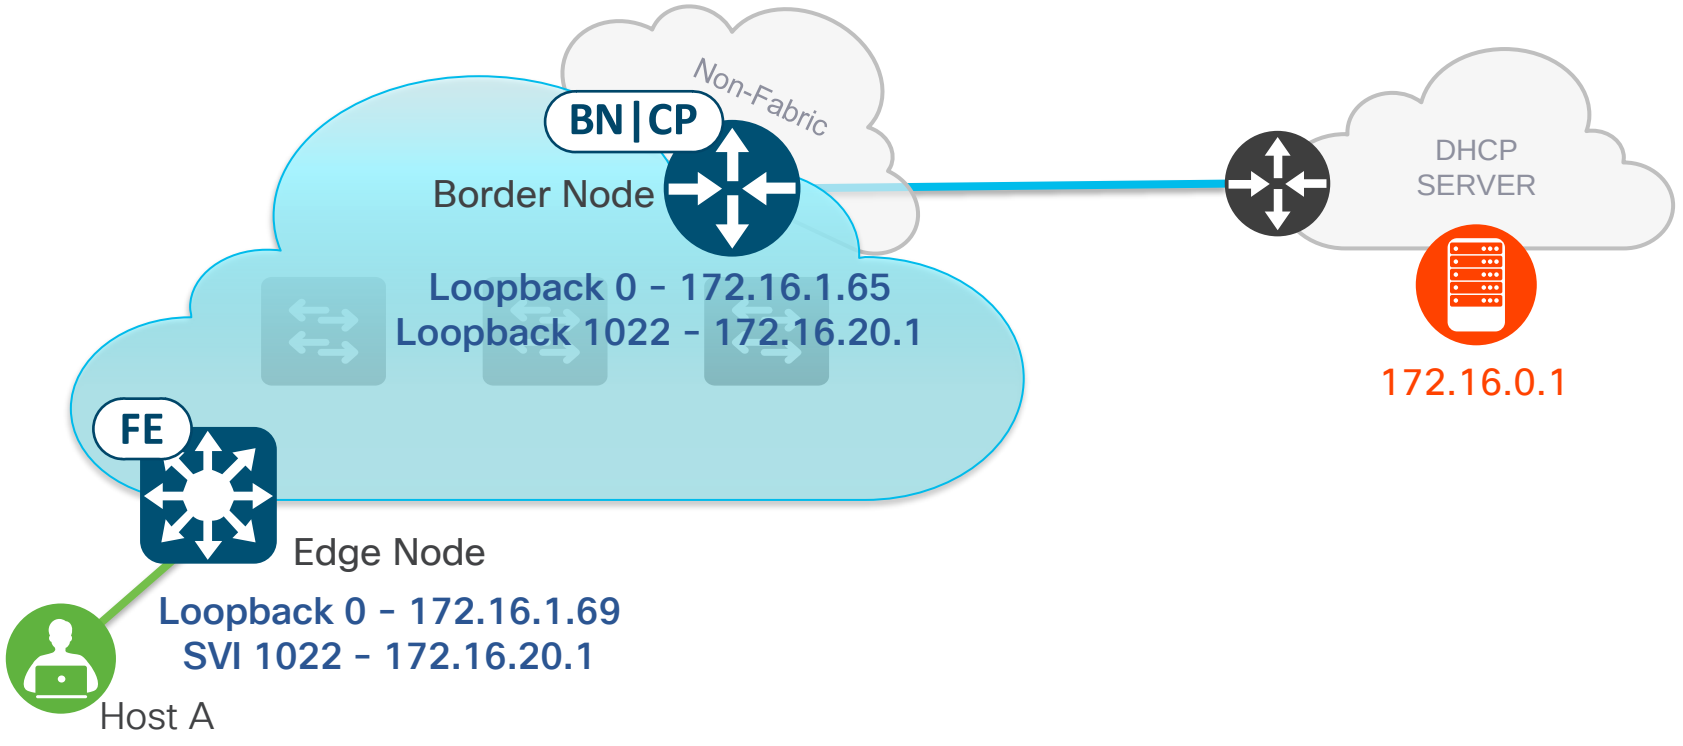

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

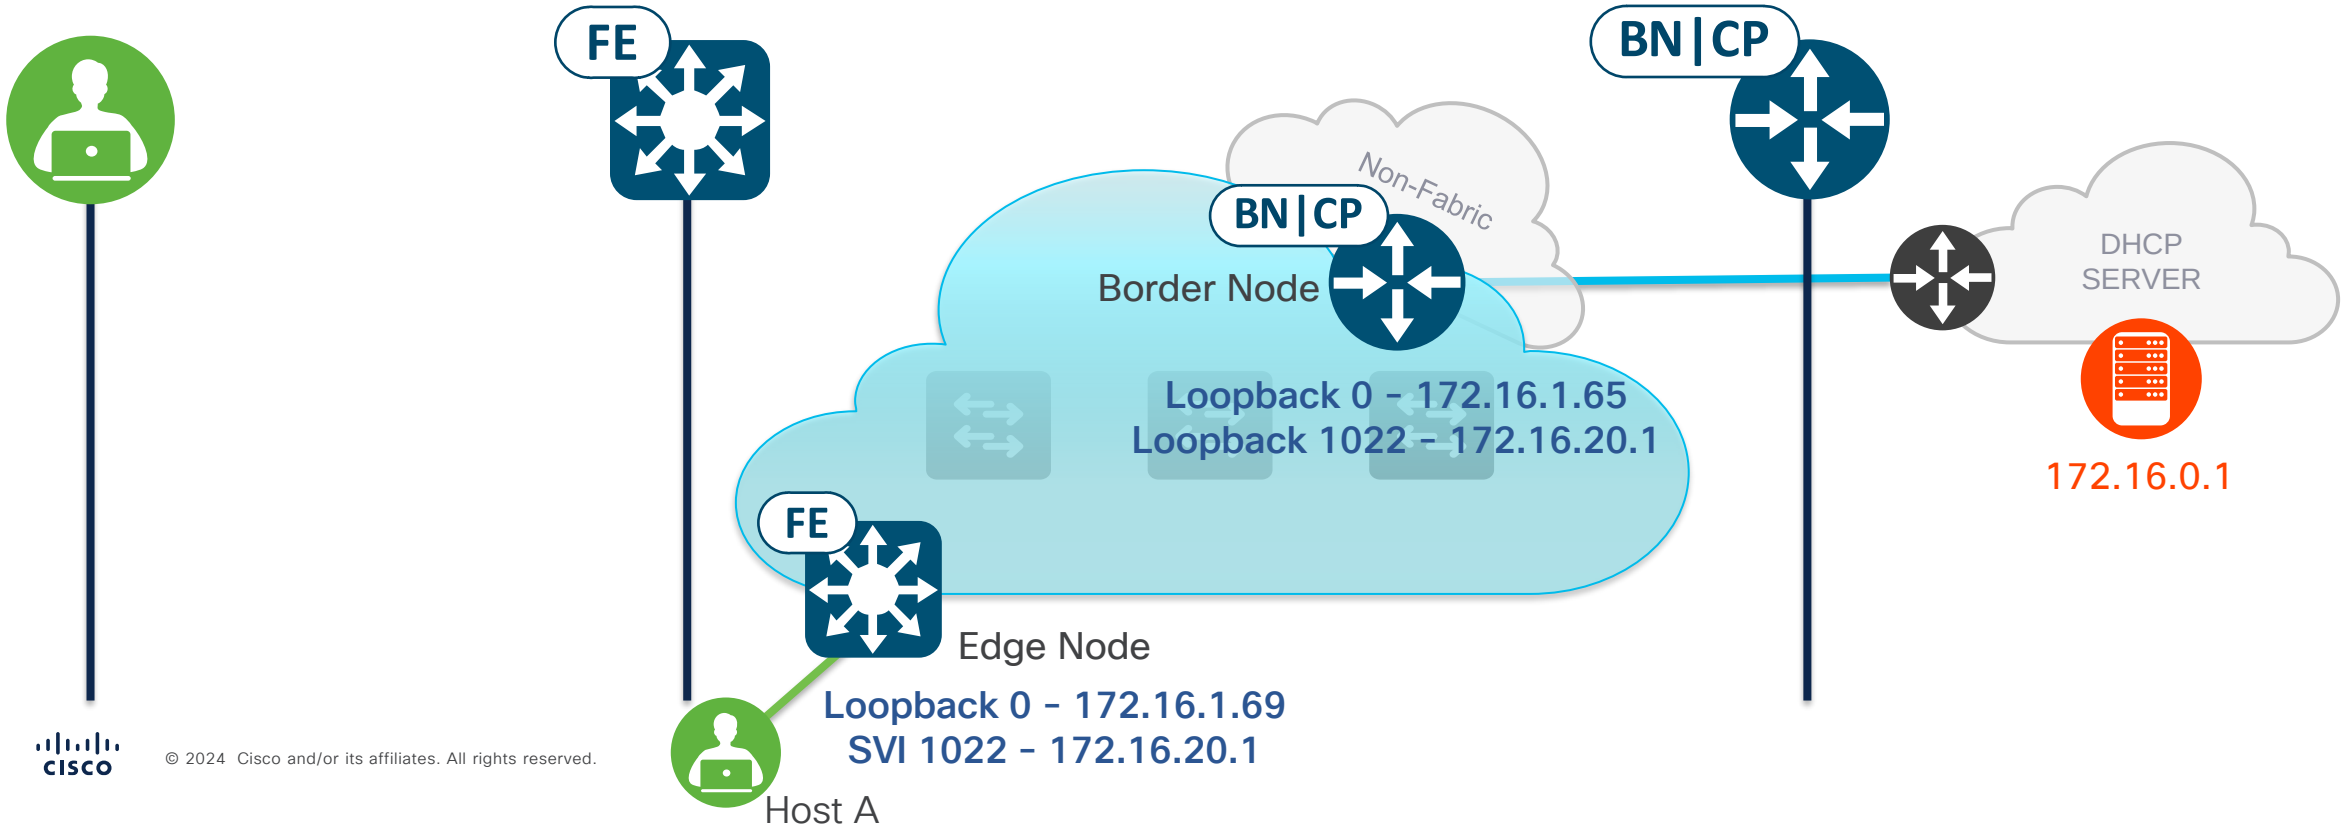

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

```
Edge1_Oeiras#show mac address interface Te1/0/3
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
1022    aaaa.bbbb.2222    DYNAMIC     Te1/0/3
```

```
Edge1_Oeiras#show ip dhcp snooping
Switch DHCP snooping is enabled
Switch DHCP gleaning is disabled
DHCP snooping is configured on following VLANs:
11,101,1021-1022,1025-1026
DHCP snooping is operational on following VLANs:
11,101,1021-1022,1025-1026

Insertion of option 82 is enabled
   circuit-id default format: vlan-mod-port
   remote-id: 00bf.7798.4980 (MAC)
```

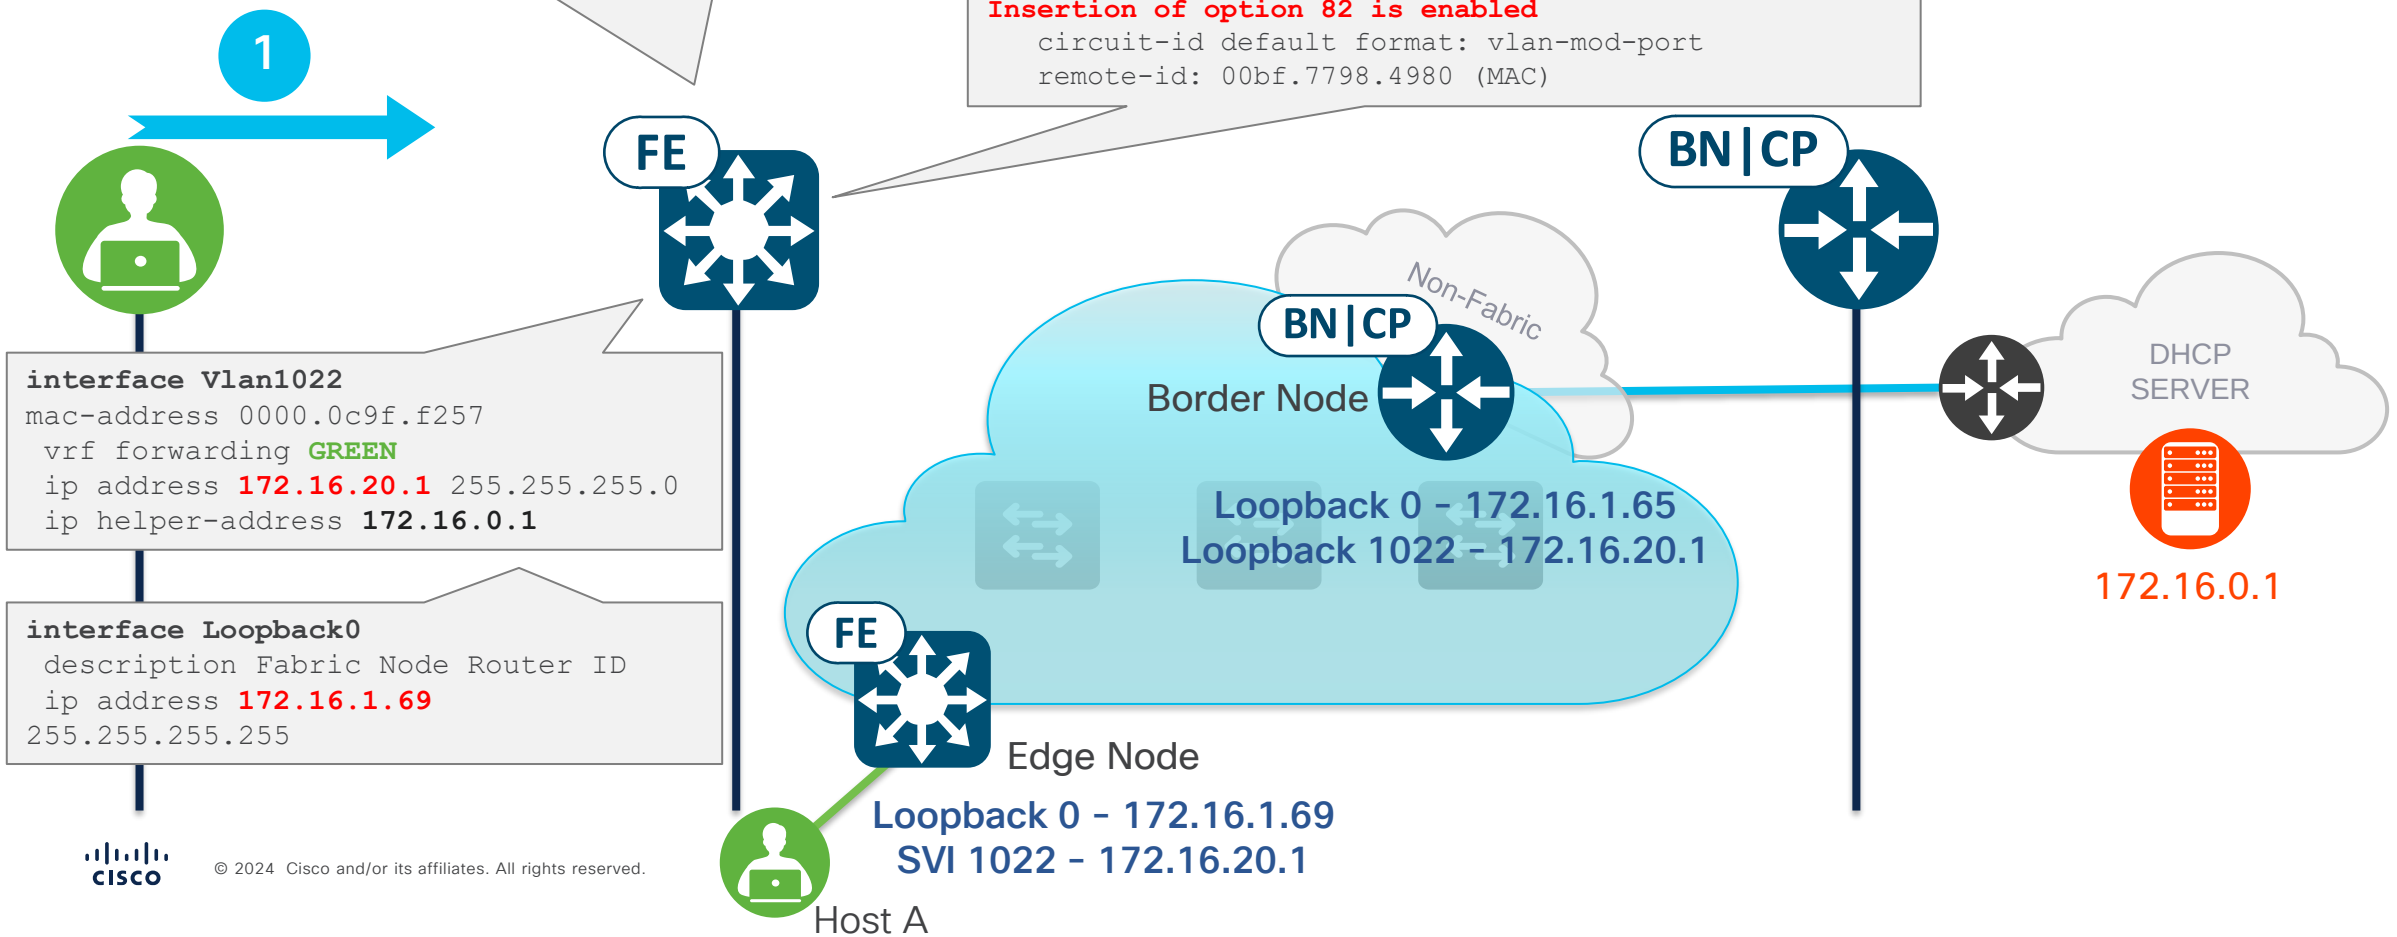

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

```
050949: *Sep  3 08:41:53.625: DHCP_SNOOPING: process new DHCP packet, message type: DHCPDISCOVER, input interface: Te1/0/3, MAC da:
ffff.ffff.ffff, MAC sa: aaaa.bbbb.2222, IP da: 255.255.255.255, IP sa: 0.0.0.0, DHCP ciaddr: 0.0.0.0, DHCP yiaddr: 0.0.0.0, DHCP siaddr: 0.0.0.0,
DHCP giaddr: 0.0.0.0, DHCP chaddr: aaaa.bbbb.2222, efp_id: 696189056, vlan_id: 1022, bootpflag:0x32768(Broadcast)
050950: *Sep  3 08:41:53.625: DHCP_SNOOPING: add relay information option.
050951: *Sep  3 08:41:53.625: DHCP_SNOOPING: Encoding opt82 CID in vlan-mod-port format
050952: *Sep  3 08:41:53.625: :VLAN case : VLAN ID 1022
050953: *Sep  3 08:41:53.625:  VRF id is valid
050954: *Sep  3 08:41:53.625: LISP ID is valid, encoding RID in srloc format
050955: *Sep  3 08:41:53.625: DHCP_SNOOPING: binary dump of relay info option, length: 22 data:
050956: *Sep  3 08:41:53.626: DHCPS BRIDGE PAK: vlan=1022 platform_flags=1
050957: *Sep  3 08:41:53.626: DHCP_SNOOPING: bridge packet get invalid mat entry: FFFF.FFFF.FFFF, packet is flooded to ingress VLAN: (1022)
050958: *Sep  3 08:41:53.626: DHCP_SNOOPING: bridge packet send packet to cpu port: Vlan1022.
```

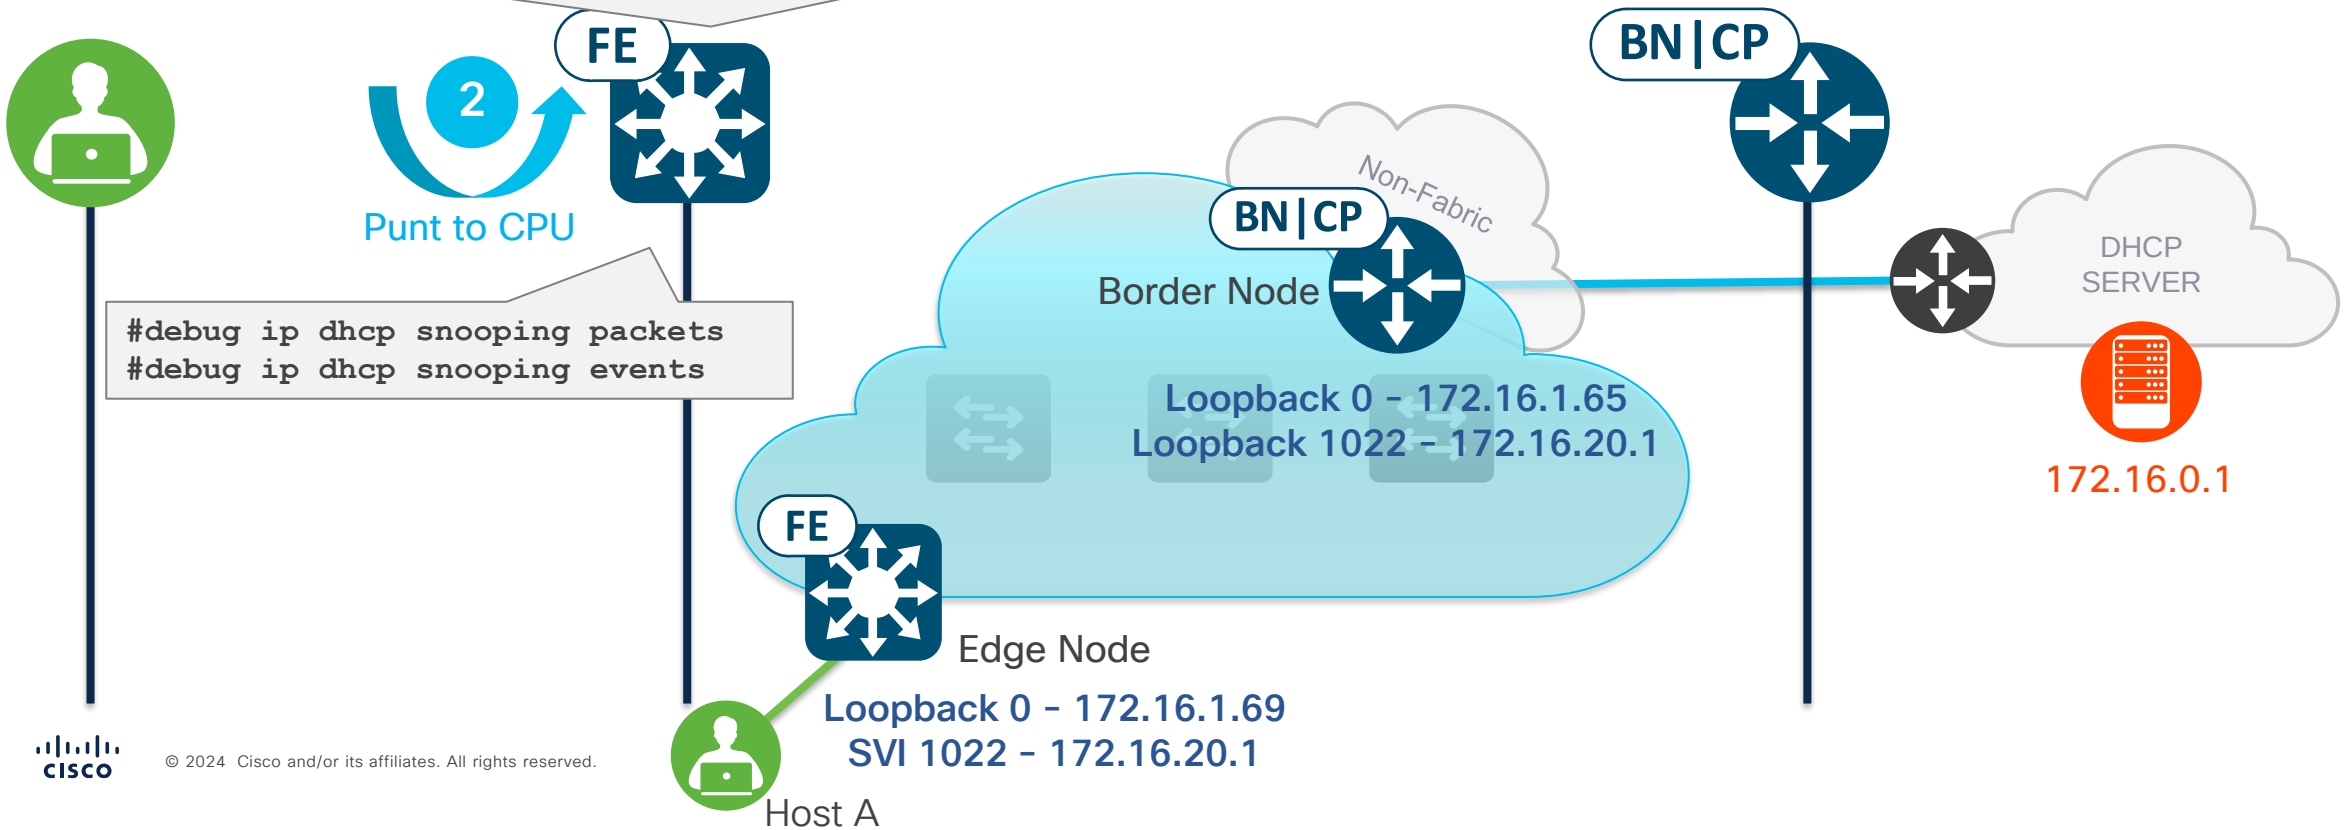

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access - Discover

Capture in CPU- Edge Node

```
Edge1_Oeiras#show monitor capture cap buffer display-filter bootp

   94  0.0.0.0 -> 255.255.255.255 DHCP 378 DHCP Discover - Original
   95  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Relay Injected
```

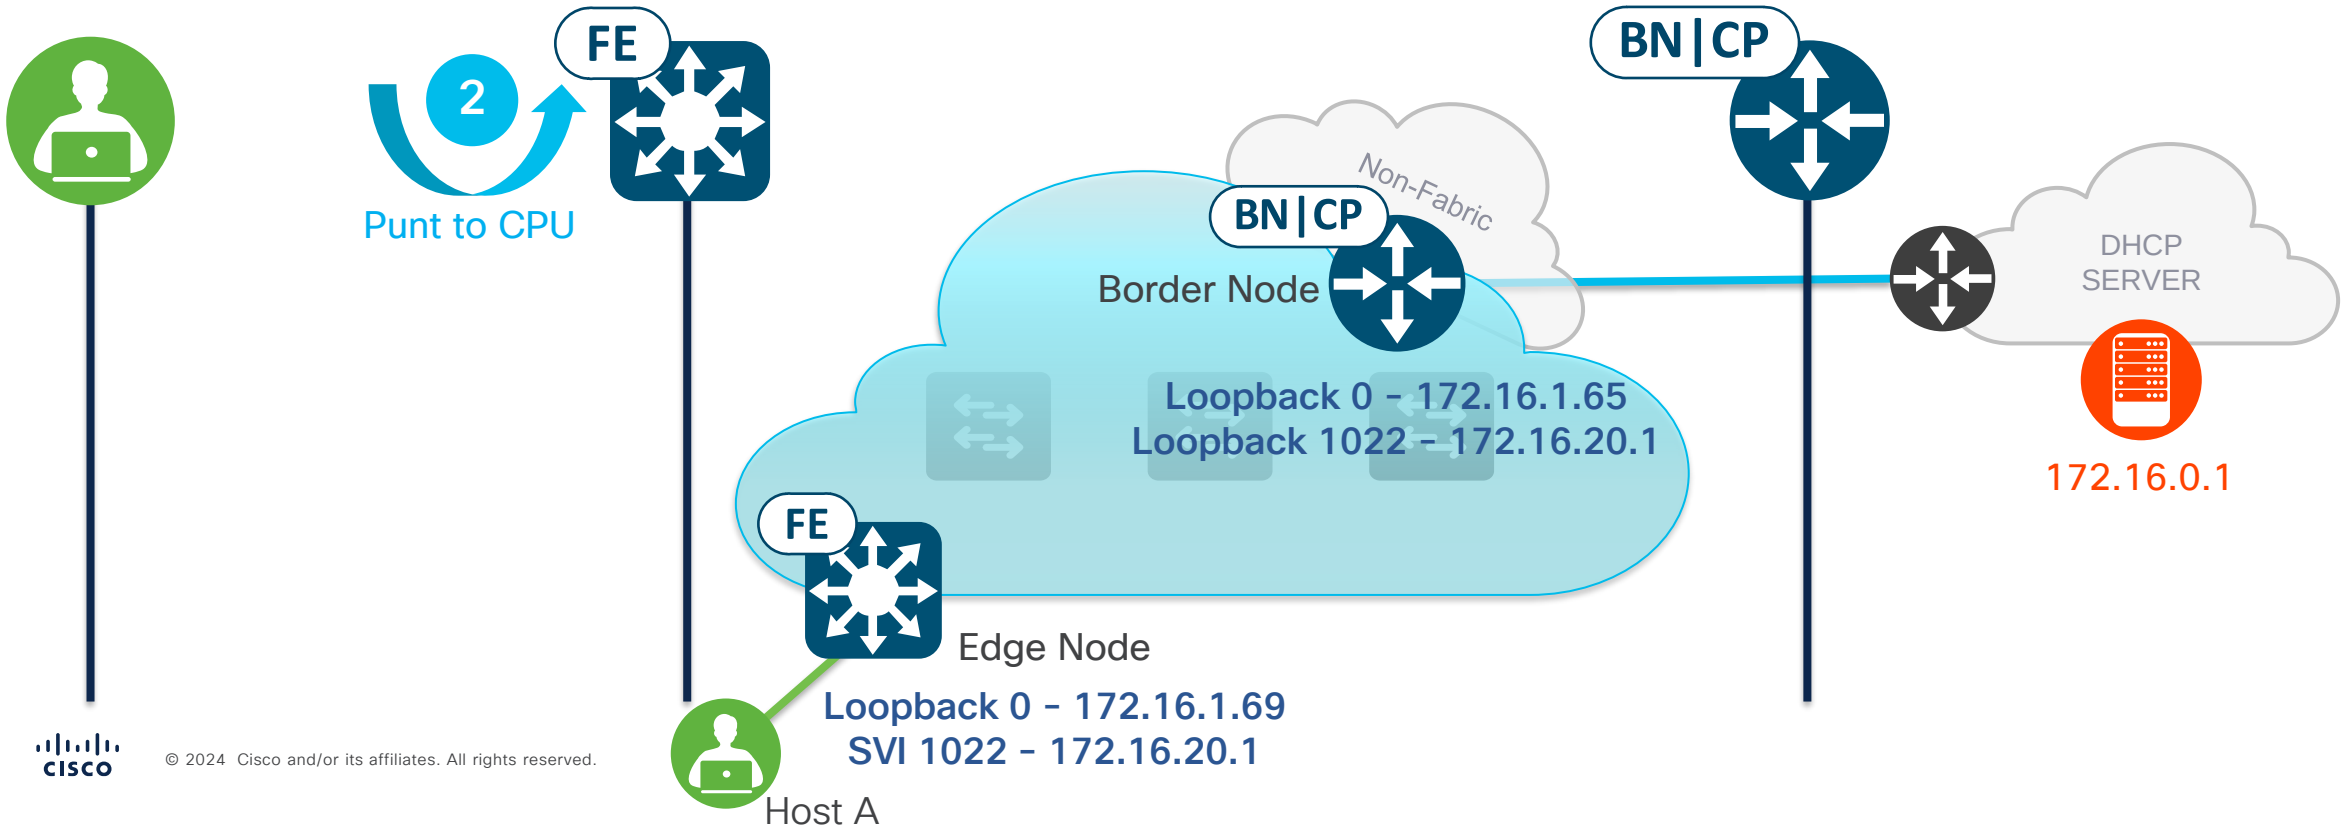

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

Capture in CPU- Edge Node

- Option 82 Inserts the Loopback 0 IP of the RLOC in Hexadecimal

ac - 172
10 - 16
01 - 1
45 - 69

172.16.1.69 > > >

```
Edge1_Oeiras#sh mon cap cap buf dis frame.number==95 de | se Agent
Option: (82) Agent Information Option
        Length: 20
        Option 82 Suboption: (1) Agent Circuit ID
            Length: 6
            Agent Circuit ID: 000403fe0103
        Option 82 Suboption: (2) Agent Remote ID
            Length: 10
            Agent Remote ID: 030800100601ac100145
```

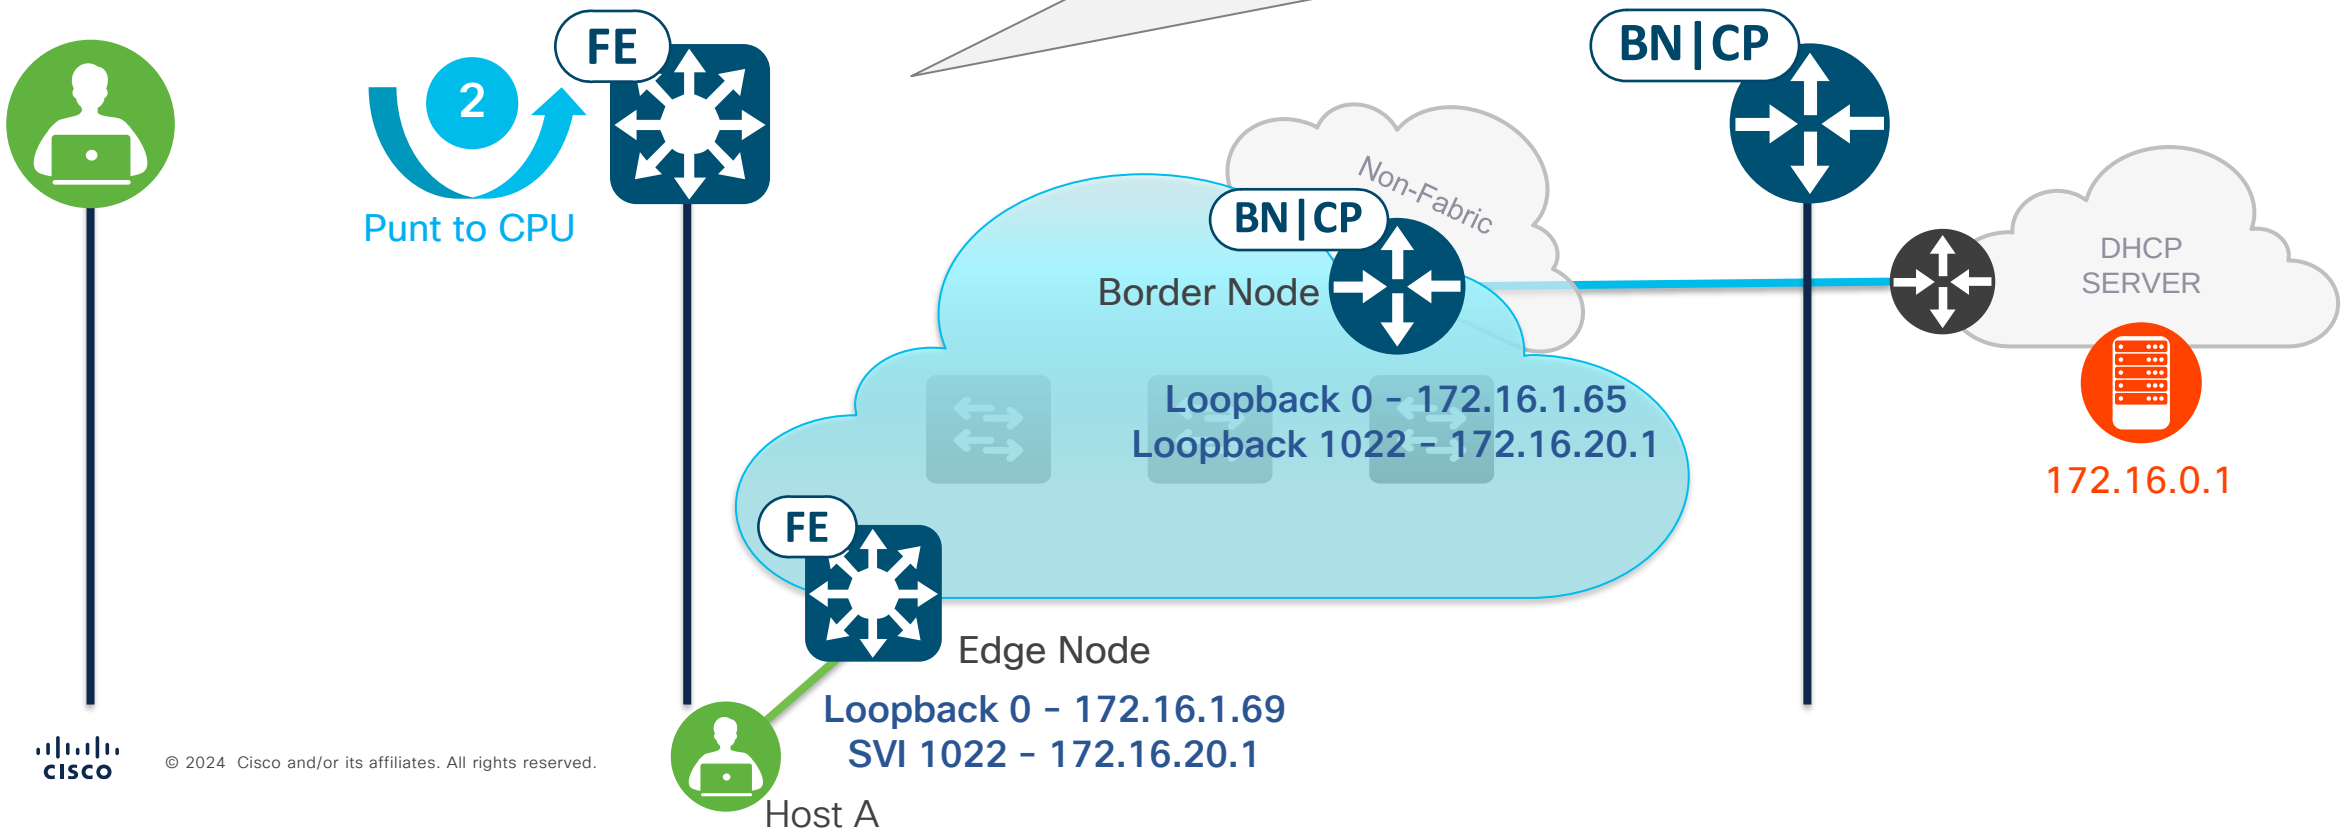

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

- The traffic is sent from 172.16.20.1 (SVI on the Edge) to 172.16.0.1 (IP of the DHCP Server)

Verifying:
- What decision does CEF make to forward the traffic?
- Is the route to the Border (RLOC) a /32?

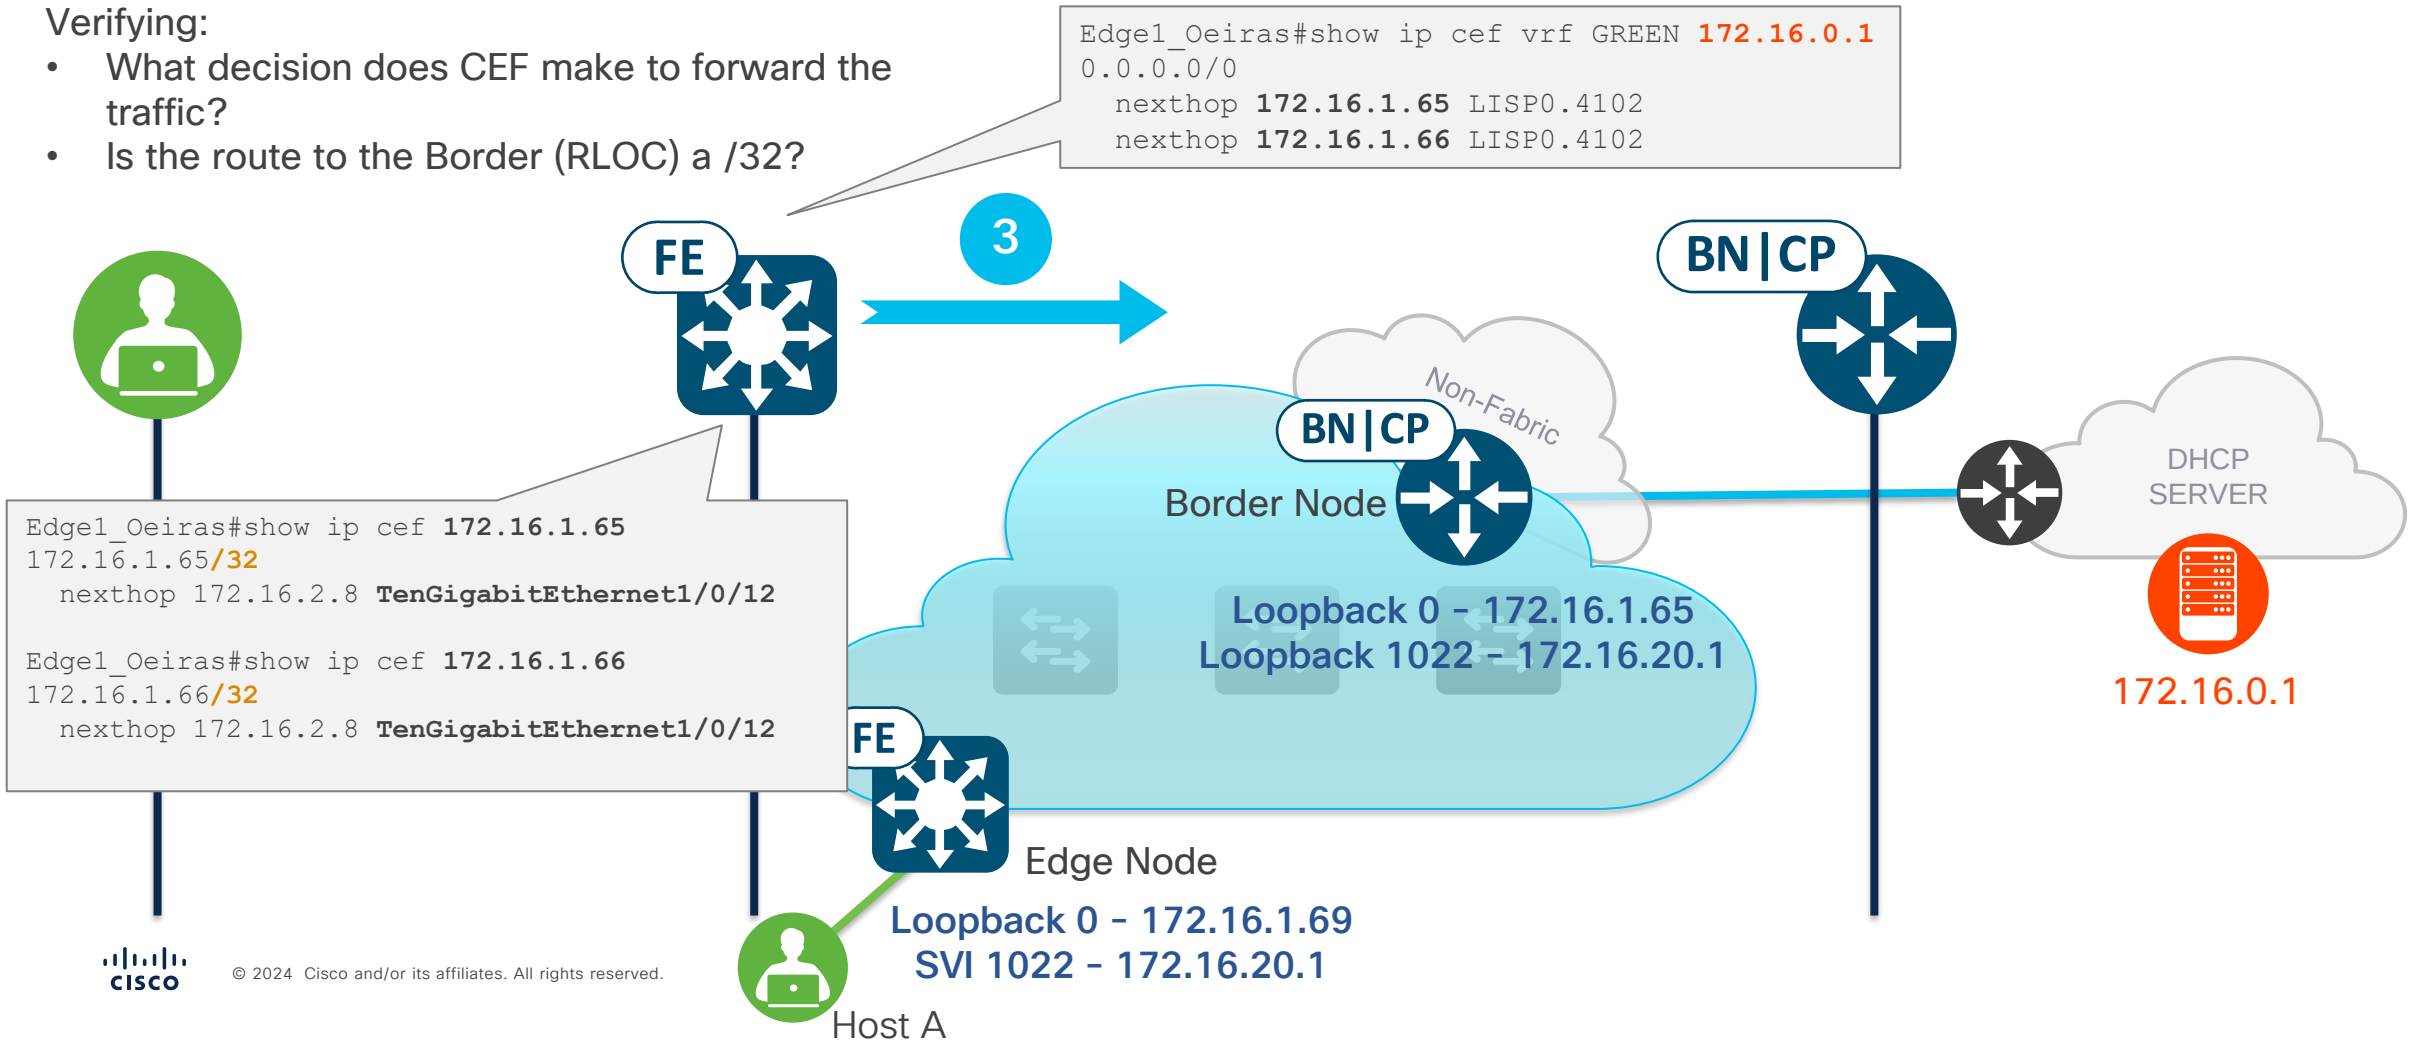

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Discover

Verifying:

- Connectivity from the Border to the DHCP Server (Routing Table)
- Correctly functioning VRF leaking on the Fusion
- Option 82 inserted in the DHCP Discover

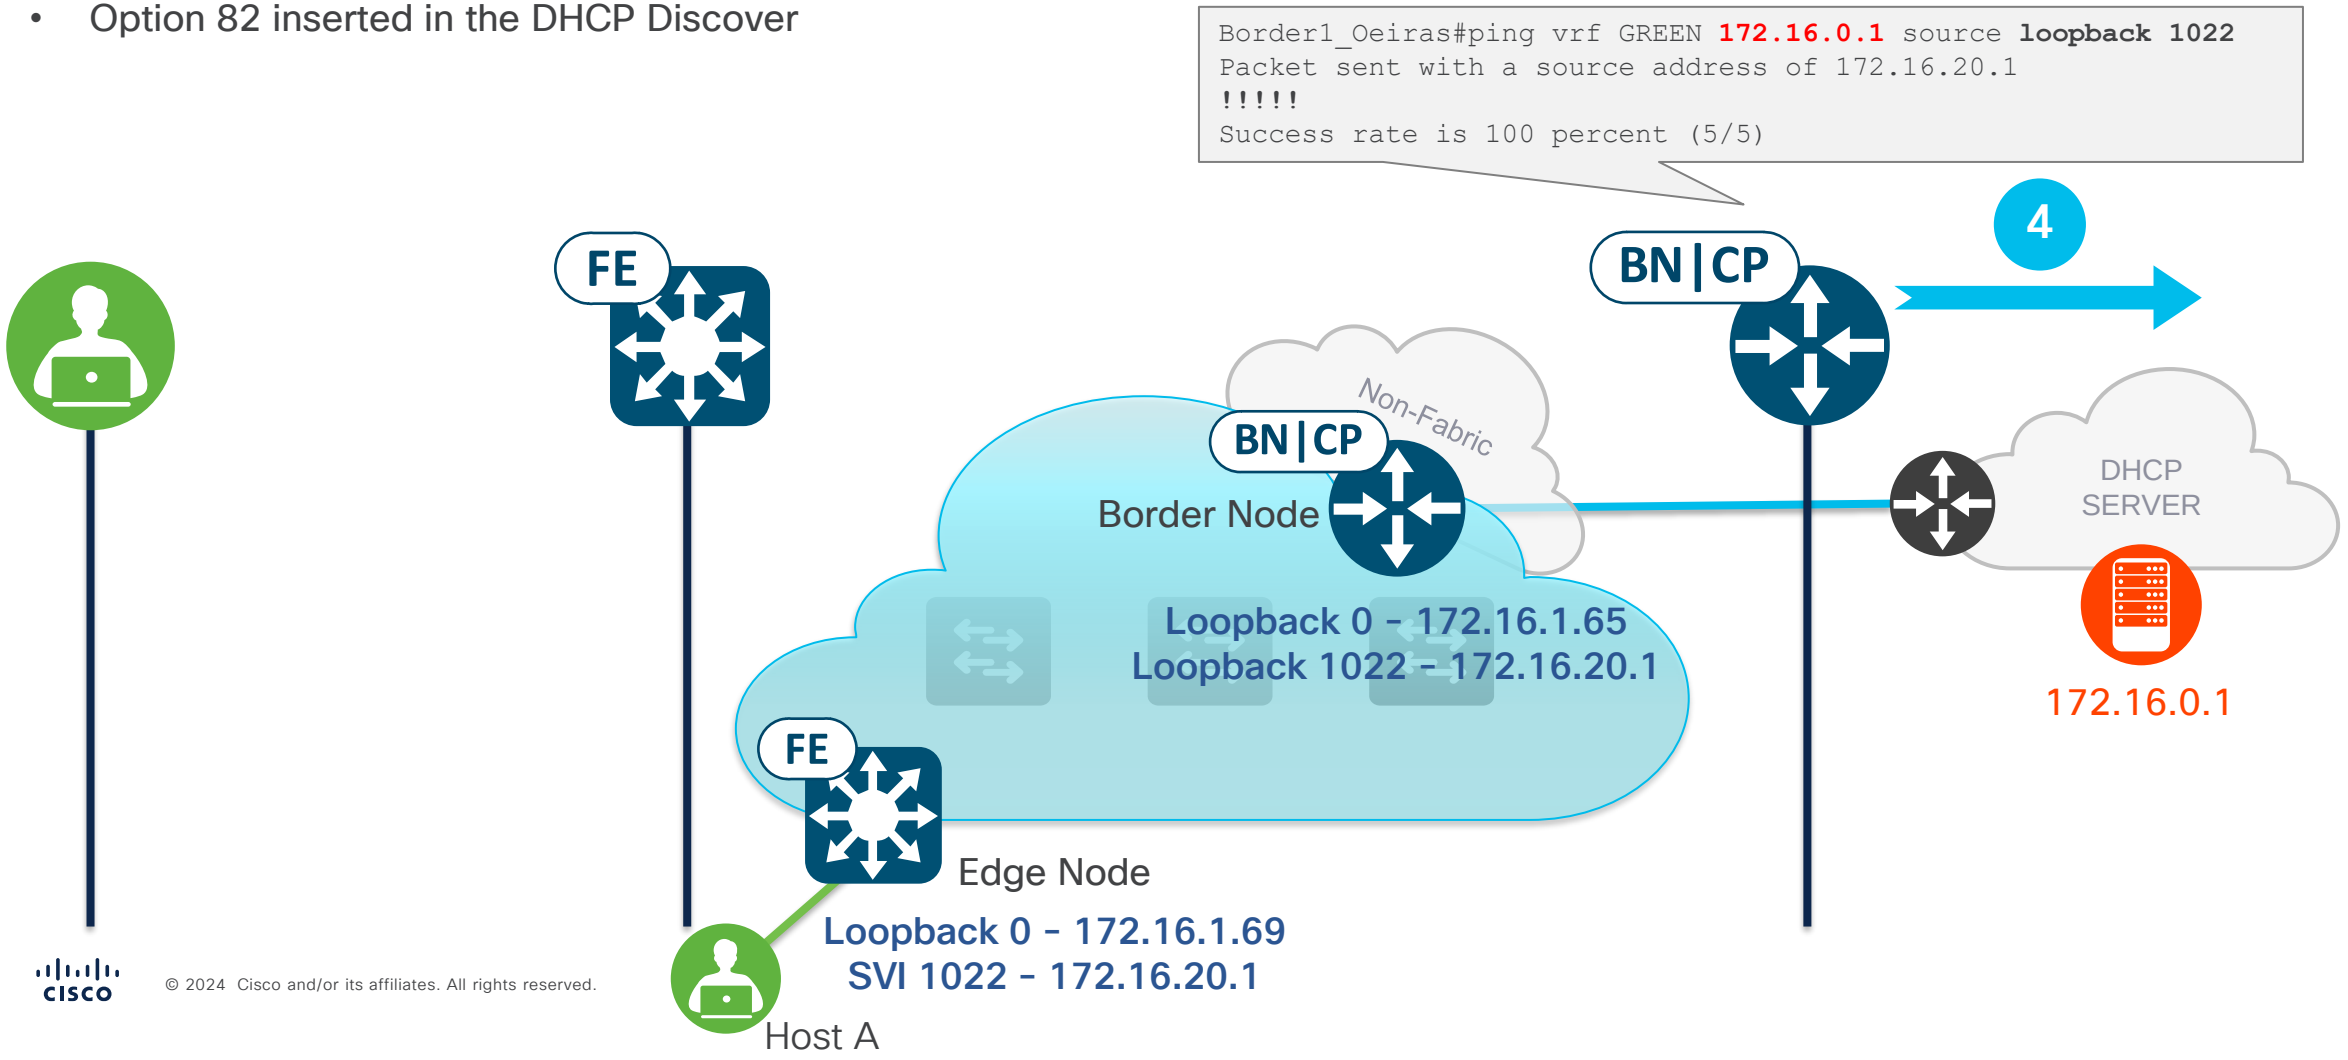

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Offer

The Offer Packet Has a Destination IP of 172.16.20.1 (Anycast Gateway)

Does Option 82 prevail in the Offer packet? –> Packet capture on the interface

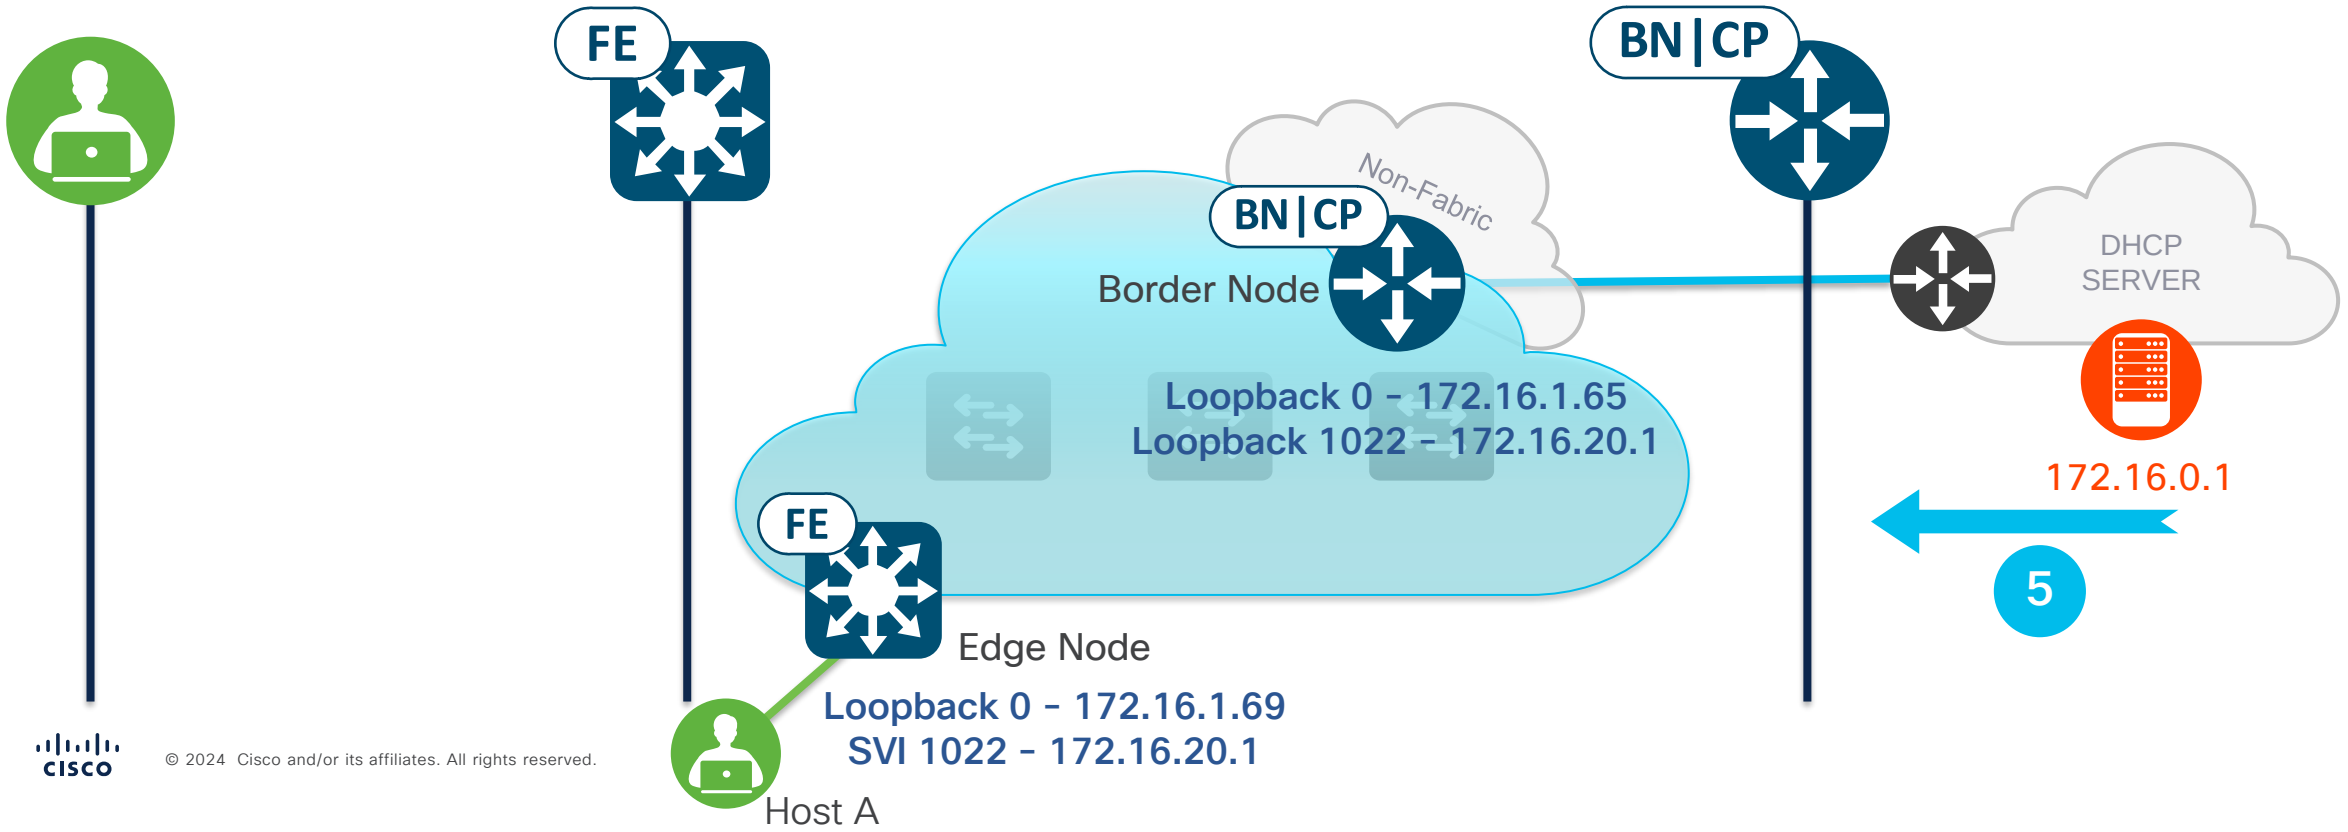

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Offer

```
Border1_Oeiras#show ip dhcp snooping
Switch DHCP snooping is disabled
Switch DHCP gleaning is disabled
DHCP snooping is configured on following VLANs:
None
```

```
Border1_Oeiras#show ip cef vrf GREEN 172.16.20.1
172.16.20.1/32
  receive for Loopback1022
```

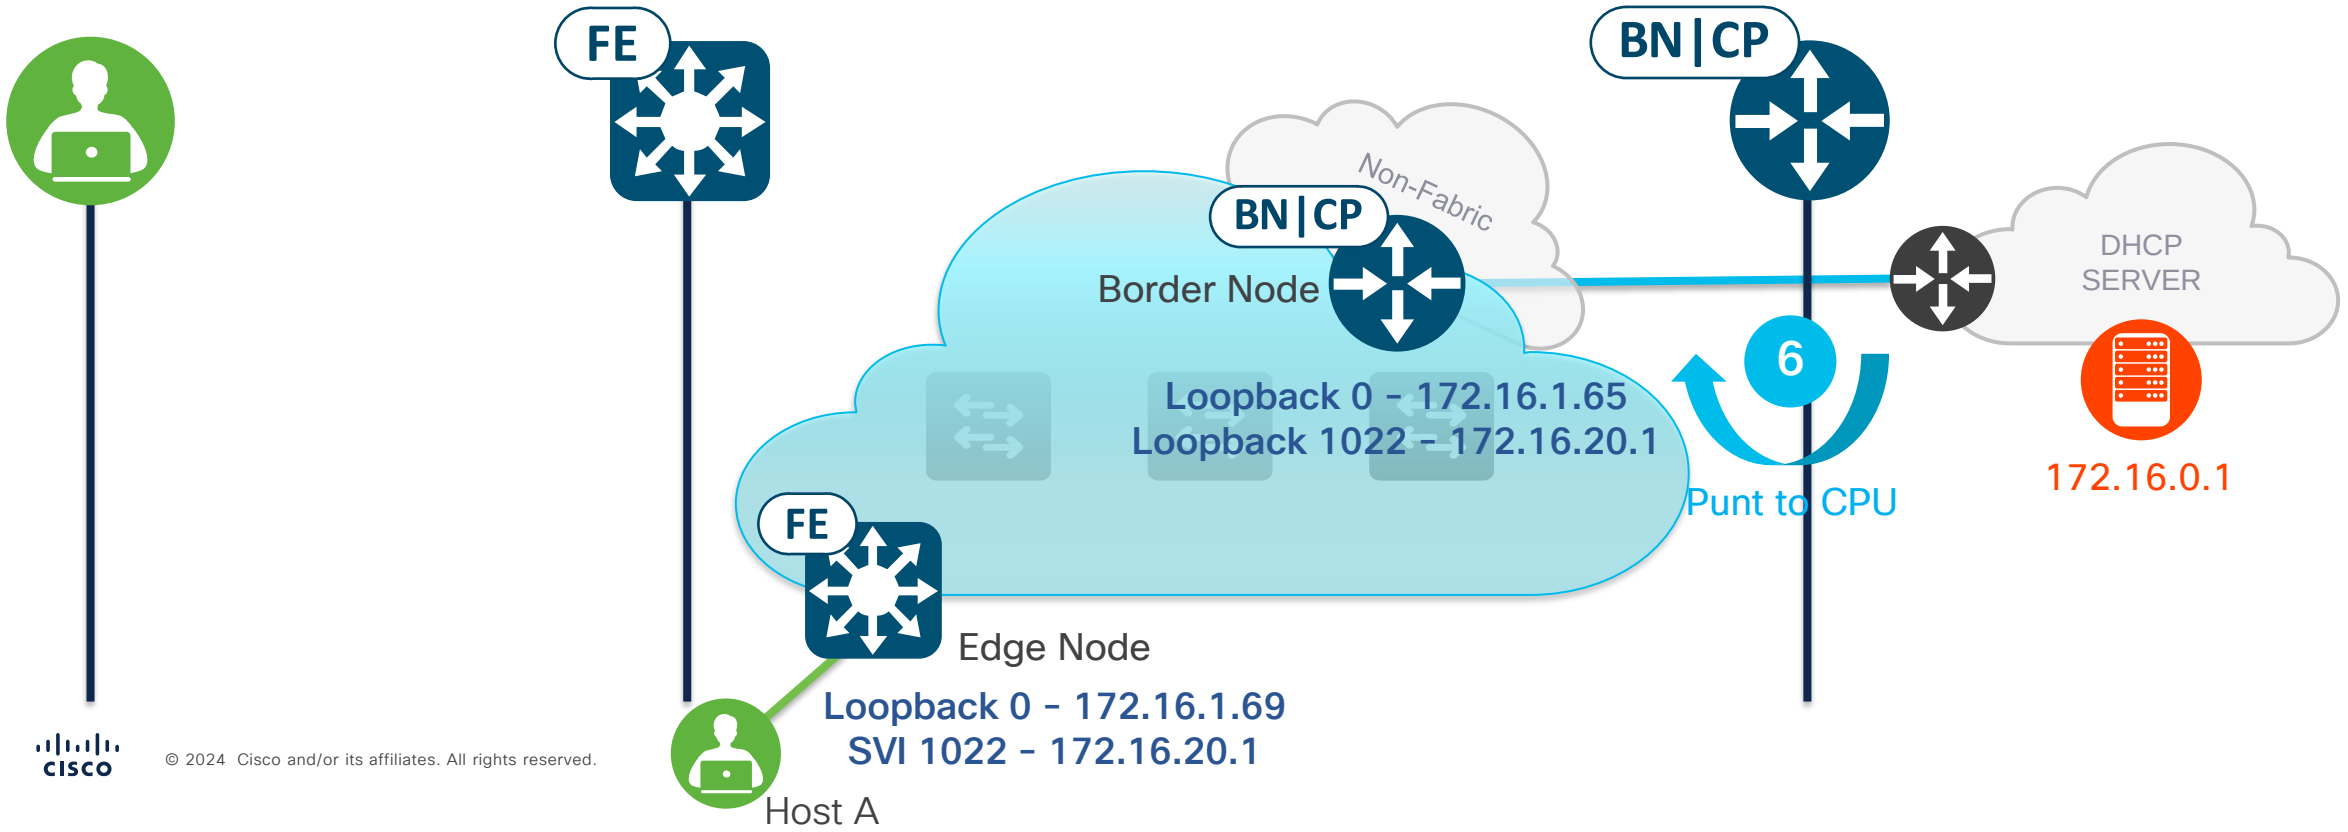

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Offer

Capture in CPU – Border Node

```
Border1_Oeiras#show monitor capture cap buffer display bootp

  166  172.16.0.1 -> 172.16.20.1  DHCP 385 DHCP Offer   -  Original
  167  172.16.0.1 -> 172.16.20.1  DHCP 431 DHCP Offer   -  VXLAN encapsulated
```

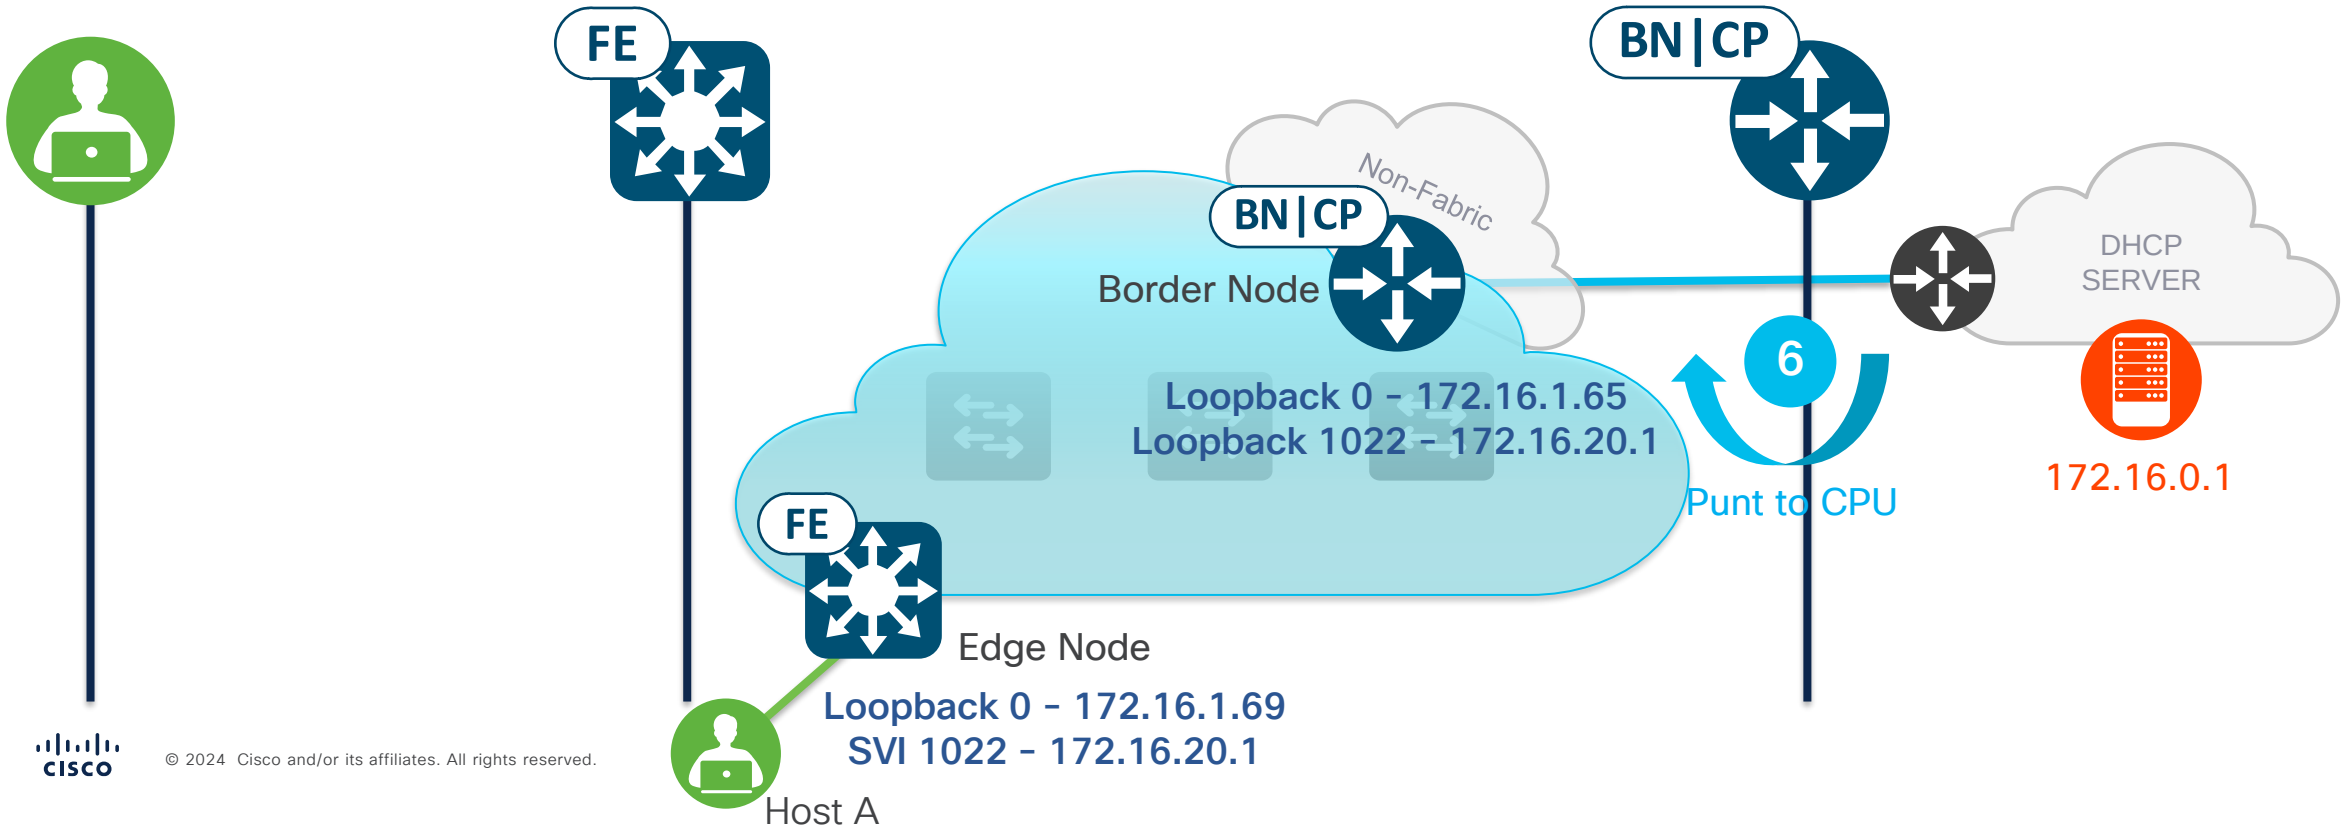

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Offer

Capture in CPU – Border Node

```
Border1_Oeiras#show monitor capture cap buffer display frame.number== 167

Internet Protocol Version 4, Src: 172.16.1.65, Dst: 172.16.1.69
    0100 .... = Version: 4
...
Internet Protocol Version 4, Src: 172.16.0.1, Dst: 172.16.20.1
    0100 .... = Version: 4
```

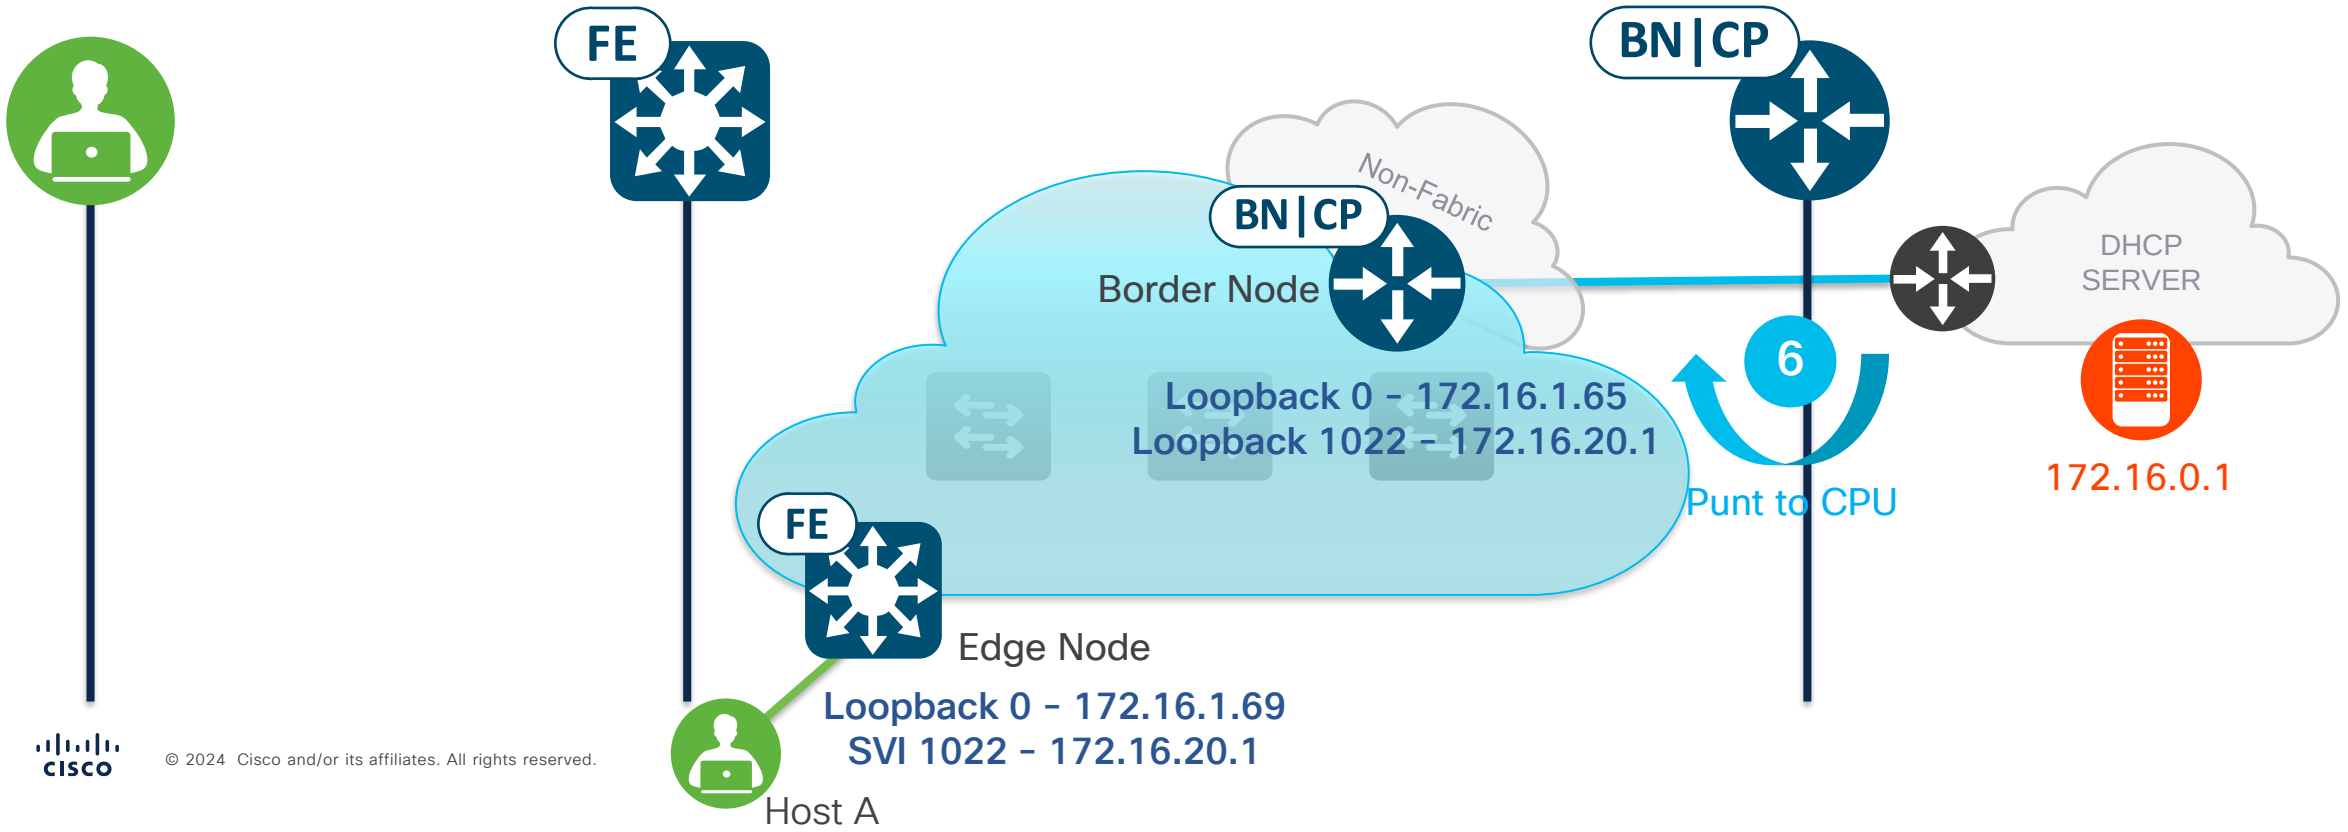

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access - Offer

Offer Packet Encapsulated in VXLAN Delivered to Edge Node
VXLAN Header Removed
The Offer Packet Originally Destined for IP 172.16.20.1 is Processed in CPU

```
Edge1_Oeiras#sh mon cap cap buf dis bootp

 280   172.16.0.1 -> 172.16.20.1  DHCP 381 DHCP Offer
```

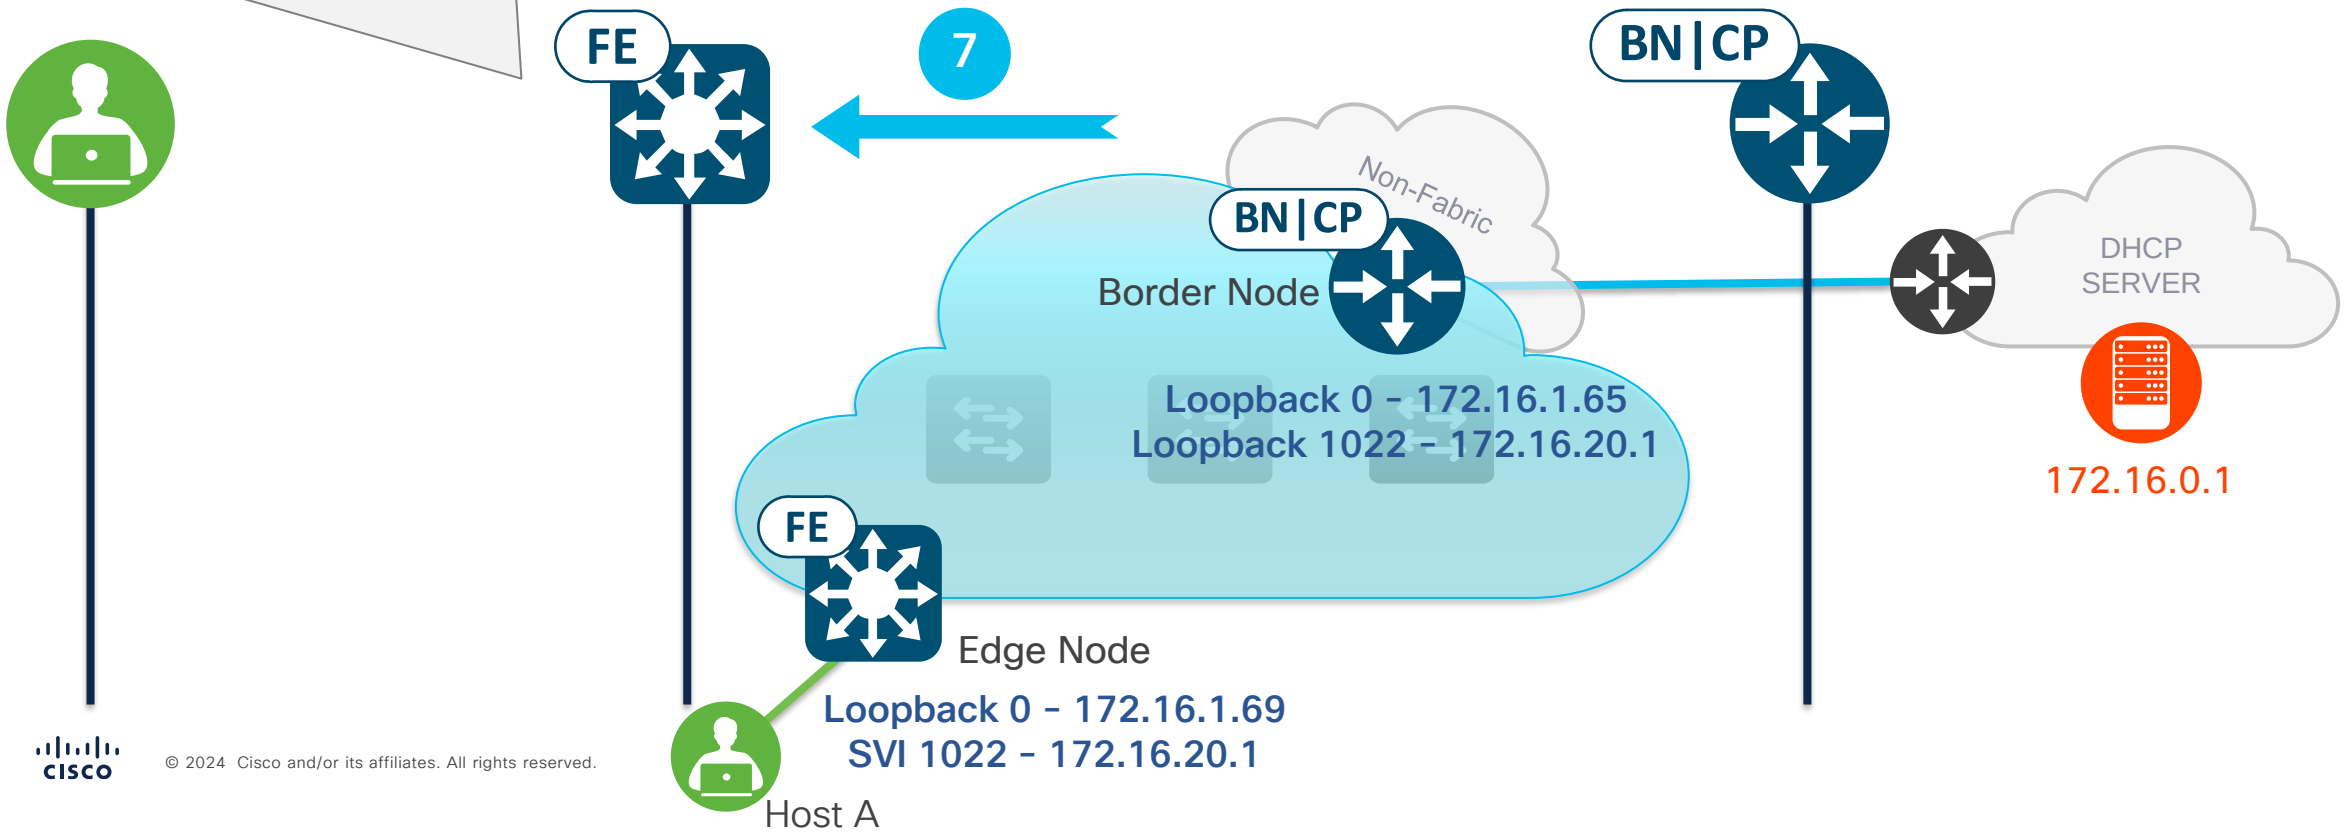

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access - Offer

Option 82 is Evaluated and Delivered to the Client

```
: *Sep  3 08:15:08.935: DHCP_SNOOPING: process new DHCP packet, message type: DHCPOFFER, input interface: Vl1022, MAC da: ffff.ffff.ffff,
MAC sa: 0000.0c9f.f257, IP da: 255.255.255.255, IP sa: 172.16.20.1, DHCP ciaddr: 0.0.0.0, DHCP yiaddr: 172.16.20.2, DHCP siaddr: 0.0.0.0,
DHCP giaddr: 172.16.20.1, DHCP chaddr: aaaa.bbbb.2222, efp_id: 696189056, vlan_id: 1022, bootpflag:0x32768(Broadcast)
: *Sep  3 08:15:08.935: DHCP_SNOOPING: binary dump of option 82, length: 22 data:
...
: *Sep  3 08:15:08.935: DHCP_SNOOPING: opt82 data indicates local packet
: *Sep  3 08:15:08.935: DHCP_SNOOPING: remove relay information option.
: *Sep  3 08:15:08.948: DHCP_SNOOPING: direct forward dhcp replyto output port: TenGigabitEthernet1/0/3.
```

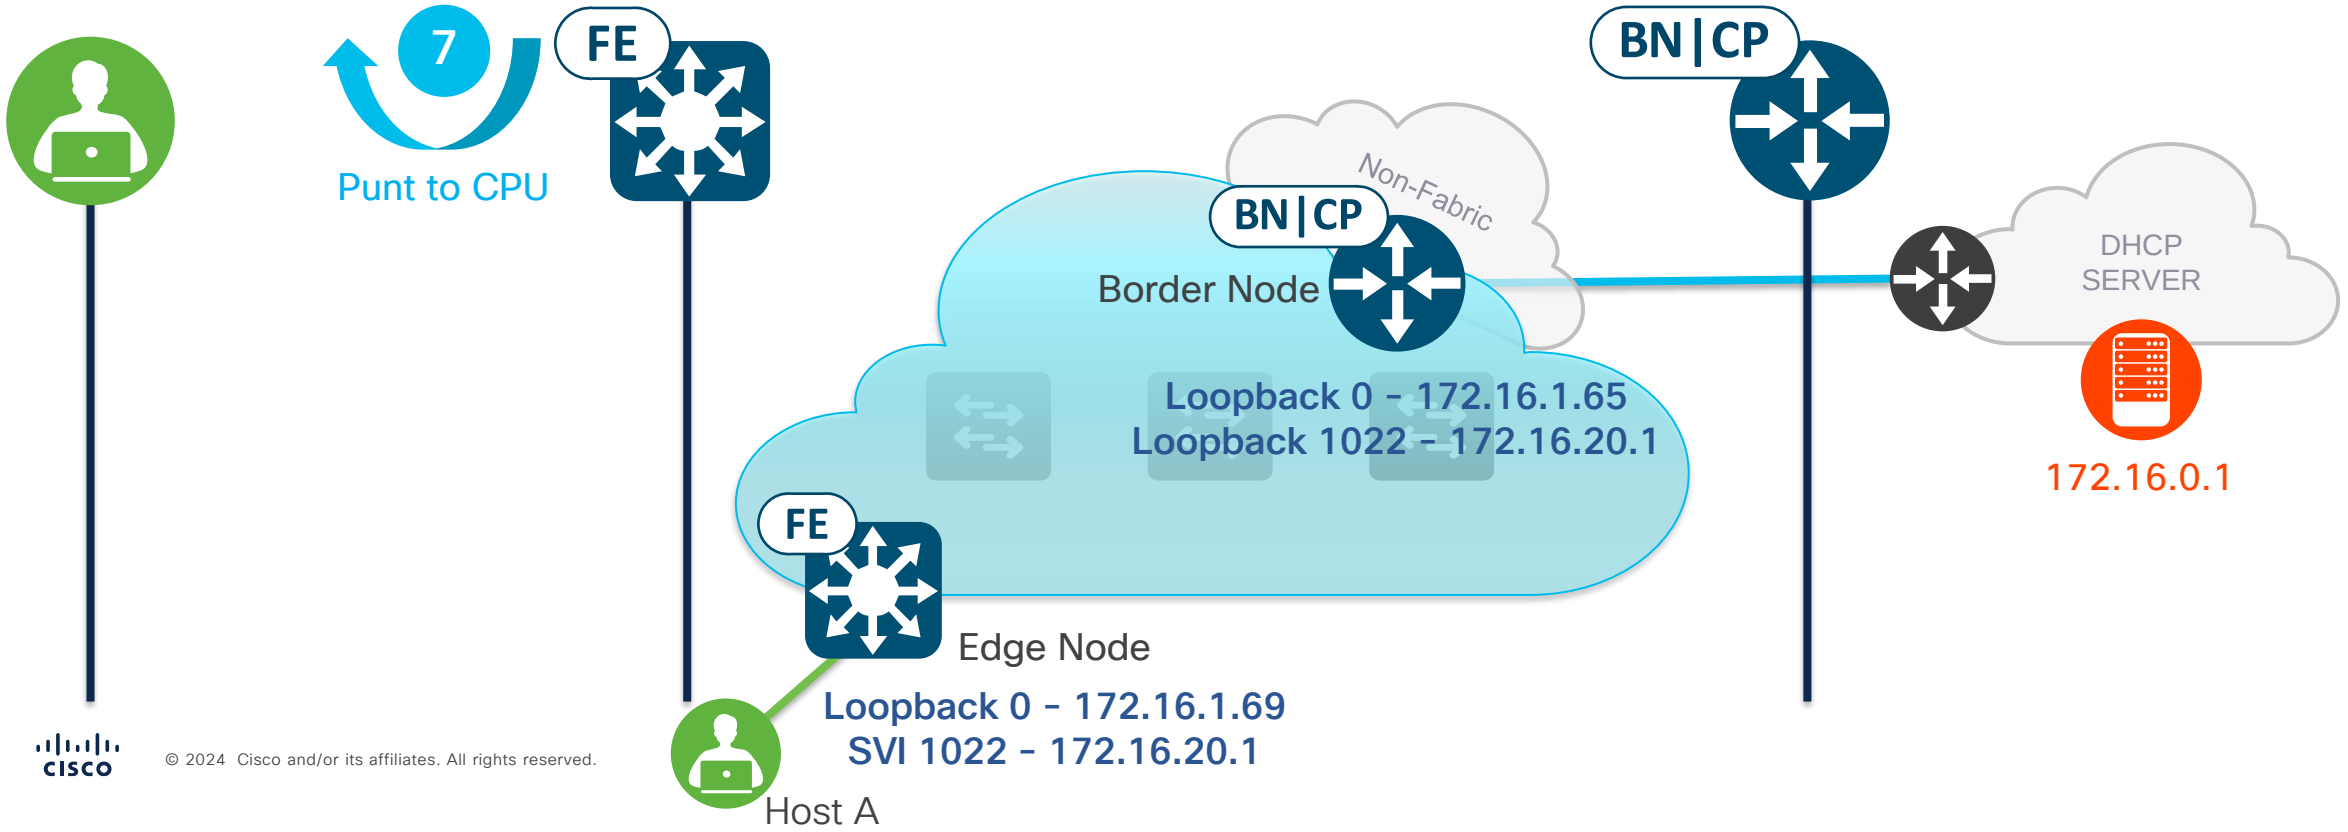

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Req / Ack

- The DHCP Request message follows exactly the same path and logic as the DHCP Discover message.
- The DHCP Acknowledge message follows exactly the same path and logic as the DHCP Offer message.

- After the DHCP Acknowledge message is processed by the Edge, the following entries are created:
  -> A DHCP Snooping binding entry is created on the Edge node
  -> A Device Tracking entry is created on the Edge node

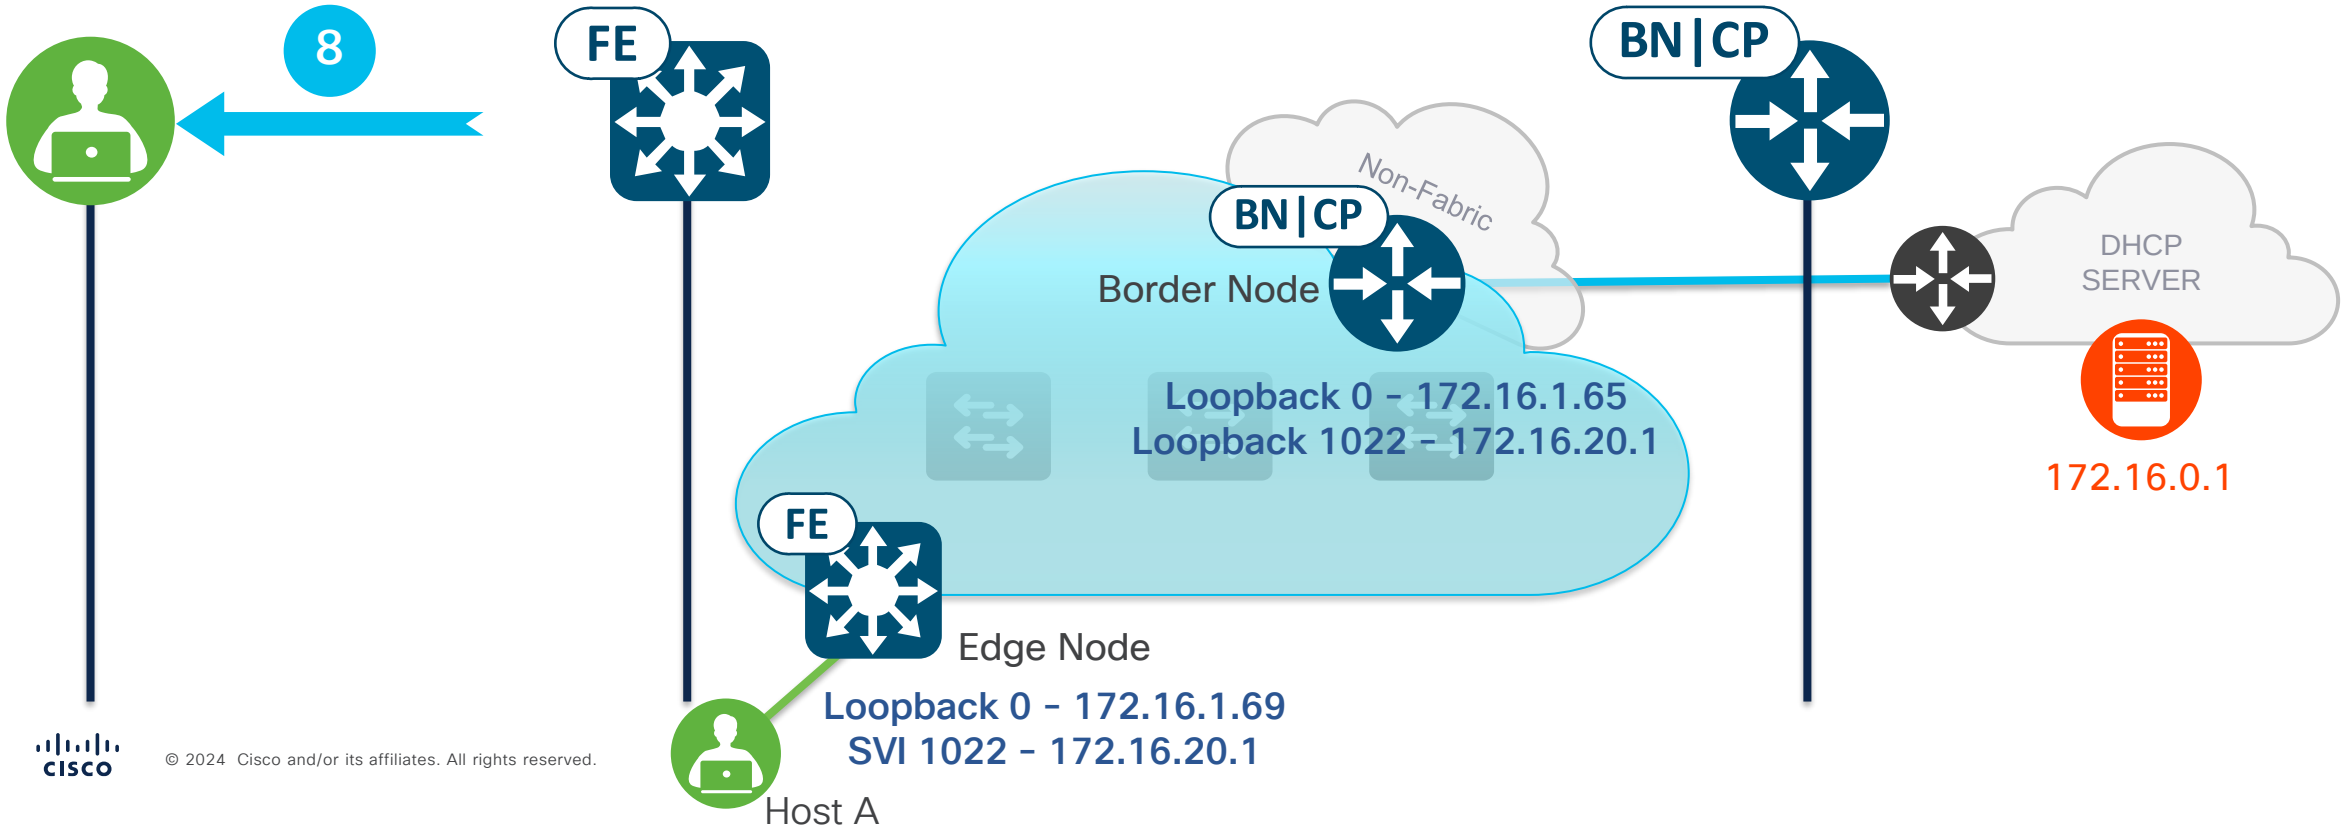

© 2024 Cisco and/or its affiliates. All rights reserved.

DHCP process within SD-Access – Req / Ack

```
Edge1_Oeiras#show device-tracking database interface te1/0/3

    Network Layer Address      Link Layer Address     Interface  vlan    prlvl      age     state      Time left
DH4 172.16.20.2                   aaaa.bbbb.2222      Te1/0/3   1022     0024      99s   REACHABLE  150 s try 0(84098 s)
```

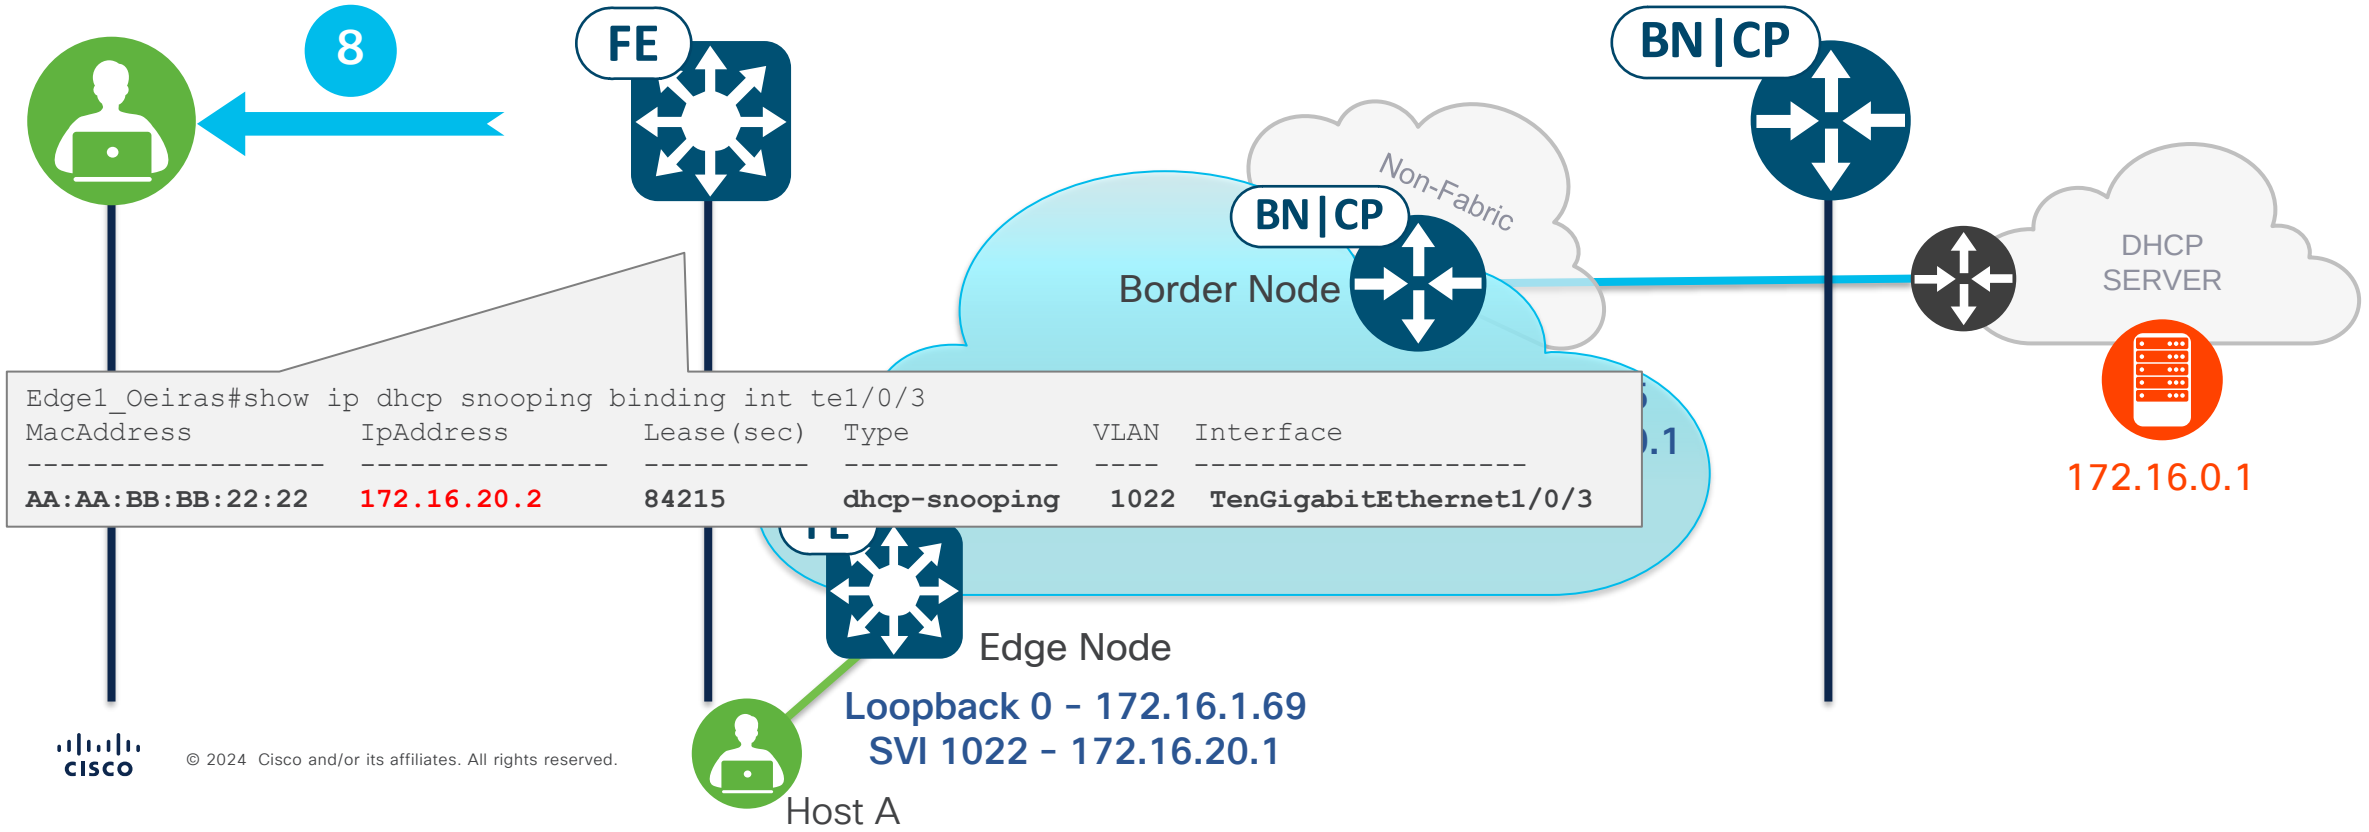


```
Edge1_Oeiras#show ip dhcp snooping binding int te1/0/3
MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
AA:AA:BB:BB:22:22   172.16.20.2      84215       dhcp-snooping   1022  TenGigabitEthernet1/0/3
```

© 2024 Cisco and/or its affiliates. All rights reserved.

Join at

slido.com

#2835 026

Passcode:

xnygs2

Which IP is inserted in option 82 so that the Border identifies the Relay Agent?

Anycast Gateway

0%

Loopback 0

0%

L3 Interface

0%

VLAN 1

0%

Troubleshooting common issues

© 2024 Cisco and/or its affiliates. All rights reserved.

Troubleshooting common issues

This section will cover three of the most common DHCP-related issues encountered in TAC cases:

1. There is no return route in the Fusion device for the relay agent’s subnet (SVI Anycast Gateway), and we don't see the Offer reaching the Border, let alone the edge node.

2. There is no route on the Border towards the DHCP server, so the Discover message will never be sent outside the Fabric.

3. DHCP snooping is disabled on the Edge node. Without DHCP snooping, we won't have Option 82, and the OFFER packet will never be forwarded from the server to the Edge node.

© 2024 Cisco and/or its affiliates. All rights reserved.

Topology

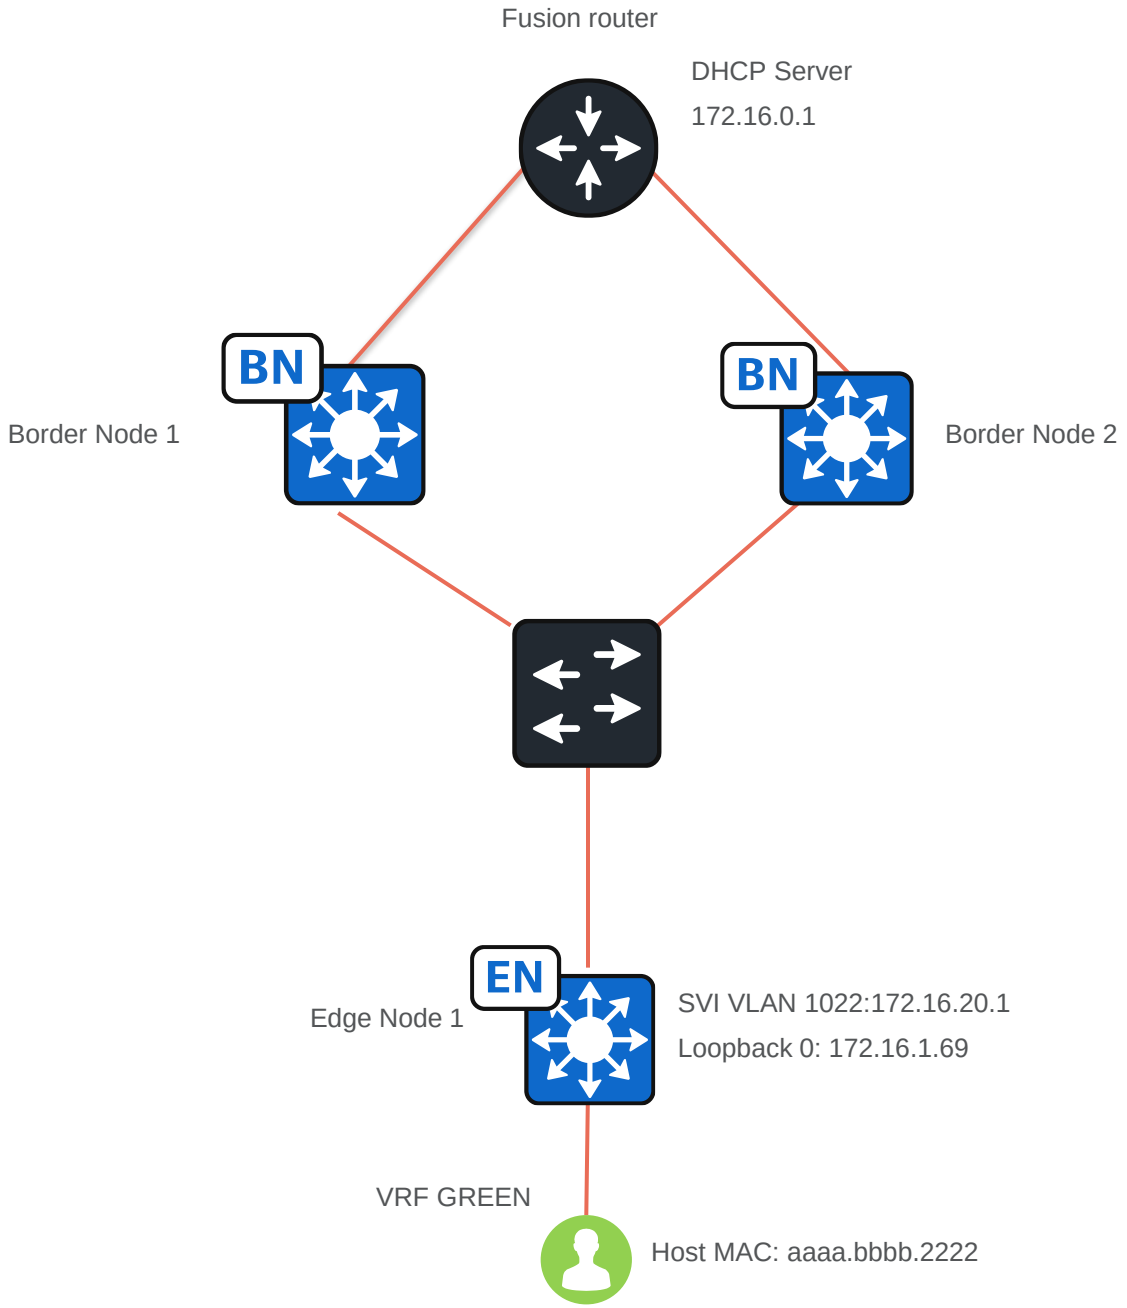

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

By capturing on the edge node in the control-plane, we can observe the complete DHCP process and isolate where the issue is occurring.

```
Edge1#monitor capture cap control-plane both match any buffer size 100 start
Edge1#show monitor capture cap buffer display-filter bootp

228   6.981436      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x864
229   6.981739      0.0.0.0 -> 255.255.255.255 DHCP 378 DHCP Discover - Transaction ID 0x864
230   6.981971  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0x864
381  10.984117      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x864
382  10.984463      0.0.0.0 -> 255.255.255.255 DHCP 378 DHCP Discover - Transaction ID 0x864
383  10.984699  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0x864
832  28.466272      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x865
833  28.466679      0.0.0.0 -> 255.255.255.255 DHCP 378 DHCP Discover - Transaction ID 0x865
834  28.466981  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0x865
```

In this scenario, the DHCP Discover message is being captured at the edge node, but the expected DHCP Offer response is missing. This indicates that the issue likely lies somewhere in the network and we can focus our troubleshooting on identifying where the OFFER packet is being lost.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

One way to review the history of DHCP packets that have been processed by DHCP snooping on the switch is by using the following command:

```
Edge1_Oeiras#show plataform dhcpsnooping client stats aaaa.bbbb.2222
…
Timestamp                Destination MAC  Destination Ip  VLAN  Message      Handler:Action
…
2024/09/02 20:27:27.531  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) PUNT:RECEIVED
2024/09/02 20:27:27.531  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) PUNT:TO_DHCPSN
2024/09/02 20:27:27.543  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) BRIDGE:RECEIVED
2024/09/02 20:27:27.543  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) BRIDGE:TO_DHCPD
2024/09/02 20:27:27.543  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) BRIDGE:TO_INJECT
2024/09/02 20:27:27.543  FFFF.FFFF.FFFF   255.255.255.255 1022 DHCPDISCOVER(B) L2INJECT:TO_FWD
2024/09/02 20:27:27.544  0000.0000.0041   172.16.0.1      0    DHCPDISCOVER(B) INJECT:RECEIVED
2024/09/02 20:27:27.544  0000.0000.0041   172.16.0.1      0    DHCPDISCOVER(B) INJECT:TO_L2FWD
…
```

Use the show clock command to check the timestamp of the packets within the output.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

Another way to check the DHCP process is by using debugs.

```
Edge1#debug ip dhcp snooping aaaa.bbbb.2222

000336: *Aug 30 18:33:54.940: DHCP_SNOOPING: process new DHCP packet, message type: DHCPDISCOVER,
input interface: Te1/0/3, MAC da: ffff.ffff.ffff, MAC sa: aaaa.bbbb.2222, IP da: 255.255.255.255, IP
sa: 0.0.0.0, DHCP ciaddr: 0.0.0.0, DHCP yiaddr: 0.0.0.0, DHCP siaddr: 0.0.0.0, DHCP giaddr: 0.0.0.0,
DHCP chaddr: aaaa.bbbb.2222, efp_id: 696189056, vlan_id: 1022, bootpflag:0x32768(Broadcast)
000337: *Aug 30 18:33:54.940: DHCP_SNOOPING: add relay information option.
…
000339: *Aug 30 18:33:54.940: DHCPS BRIDGE PAK: vlan=1022 platform_flags=1
000340: *Aug 30 18:33:54.940: DHCP_SNOOPING: bridge packet get invalid mat entry: FFFF.FFFF.FFFF,
packet is flooded to ingress VLAN: (1022)
000341: *Aug 30 18:33:54.940: DHCP_SNOOPING: bridge packet send packet to cpu port: Vlan1022
```

Just like in the capture on the control-plane, we only see the DISCOVER, but there's no sign of the OFFER. Other debugs that can provide more information about the DORA process are:

debug ip dhcp snooping events
debug ip dhcp snooping packets

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

Where will the Edge node send the DISCOVER?

```
Edge1#show ip cef vrf GREEN 172.16.0.1
172.16.0.1/32
  nexthop 172.16.1.65 LISP0.4102
  nexthop 172.16.1.66 LISP0.4102

Edge1#show lisp instance-id 4102 ipv4 map-cache 172.16.0.1
0.0.0.0/0, uptime: 12:19:19, expires: never, via static-send-map-request
  Sources: static-send-map-request
  State: send-map-request, last modified: 12:19:19, map-source: local
  Exempt, Packets out: 267(80416 bytes), counters are not accurate (~ 00:01:43 ago)
  Configured as EID address space
  Encapsulating to proxy ETR

Edge1#show run | i petr
  use-petr 172.16.1.65
  use-petr 172.16.1.66
```

How to obtain the instance-id of a VRF? Use the show vrf command, and it will be listed in the interfaces column.

```
Edge1#show vrf
  Name                             Default RD            Protocols   Interfaces
  GREEN                            <not set>             ipv4        LI0.4102
                                                                     Vl1022
```

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

Capturing on the Border, on the link towards the Fusion, we can observe the DISCOVER leaving the fabric.

```
Border1#monitor capture cap int TwentyFiveGigE1/0/2 both match any buffer size 100 start

Border1#sh monitor capture cap buff display-filter "bootp and dhcp.hw.mac_addr==aaaa.bbbb.2222"

  112   9.256495  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
  620  12.722132  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
  932  16.726726  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
…
```

Where is the OFFER?

```
Border2#show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.bbbb.2222"
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

…
```

Keep in mind that the DISCOVER packet may leave through one Border, and the OFFER may return through a different Border.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 1: No route in the Fusion

Capturing on the Fusion to confirm that the Border is sending us the DHCP Discover packet.

```
Fusion#monitor capture cap int Gig 1/0/4 both match any buff size 100 start

Fusion#sh monitor capture cap buff display-filter "bootp and dhcp.hw.mac_addr==aaaa.bbbb.2222"

  112   9.256495  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
  620  12.722132  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
  932  16.726726  172.16.20.1 -> 172.16.0.1   DHCP 428 DHCP Discover - Transaction ID 0xe24
```

Checking the route on the Fusion:

```
Fusion#show ip cef vrf GREEN 172.16.20.0
0.0.0.0/0
  no route ← No route available to reach the vlan 1022 subnet
```

Even in a scenario where we have a return route on the Fusion, is it the next hop we expect to see?

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 2: There is no route on the Border to the DHCP server.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 2: There is no route on the Border to the DHCP server.

The Border must have a route to the DHCP server. First, verify that the DISCOVER is reaching one of the Borders. Next, capture traffic on Border 1.

```
Border1#show run vrf GREEN | i  bgp|update
router bgp 65001
 neighbor 172.16.0.1 update-source Vlan82

Border1#monitor capture cap interface vlan 82 both match any buffer size 100 start
Border1#sh monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.bbbb.2222"
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    3   2.148764  172.16.20.1 -> 172.16.0.1   DHCP 386 DHCP Discover - Transaction ID 0x262f
    6   5.577630  172.16.20.1 -> 172.16.0.1   DHCP 386 DHCP Discover - Transaction ID 0x262f
    9   9.578018  172.16.20.1 -> 172.16.0.1   DHCP 386 DHCP Discover - Transaction ID 0x262f
   18  27.319294  172.16.20.1 -> 172.16.0.1   DHCP 386 DHCP Discover - Transaction ID 0x2632
```

Capture on Border 2

```
Border2#sh monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.bbbb.2222"
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
```

The SVI where the capture is performed is the one used for eBGP peering with the Fusion.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 2: There is no route on the Border to the DHCP server.

Reviewing the route on the Border:

```
Border1#show ip cef vrf GREEN 172.16.0.1
0.0.0.0/0
   no route ← No route available to reach the dhcp server
```

Example of an overlapping between the DHCP server subnet and an IP pool subnet within the Fabric in VRF GREEN:

```
Border1#show ip cef vrf GREEN 172.16.0.1
172.16.0.0/24
  attached to LISP0.4102

CEF should not use LISP to resolve the DHCP server's prefix in the Borders!
```

At this point, we need to find a way for the Border to have an outbound route to the DHCP server, either statically or via BGP.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 3: DHCP snooping is disabled on the Edge node.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 3: DHCP snooping is disabled on the Edge node.

DHCP snooping must be enabled globally and on the client VLAN.

```
Edge1#show ip dhcp snooping
Switch DHCP snooping is enabled
Switch DHCP gleaning is disabled
DHCP snooping is configured on following VLANs:
11,101,1021-1022,1025-1026
DHCP snooping is operational on following VLANs:
11,101,1021-1022,1025-1026
 …
```

If DHCP snooping is disabled, the DHCP snooping debugs will not show anything!

```
Edge1#monitor capture cap control-plane both match any buffer size 100 start
Edge1#show monitor capture cap buffer display-filter bootp
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

   52   2.701410      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x11ab
   53   2.703662  172.16.20.1 -> 172.16.0.1   DHCP 416 DHCP Discover - Transaction ID 0x11ab
  533   6.688936      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x11ab
  534   6.689350  172.16.20.1 -> 172.16.0.1   DHCP 416 DHCP Discover - Transaction ID 0x11ab
  605  10.690613      0.0.0.0 -> 255.255.255.255 DHCP 356 DHCP Discover - Transaction ID 0x11ab
  606  10.691020  172.16.20.1 -> 172.16.0.1   DHCP 416 DHCP Discover - Transaction ID 0x11ab
```

An indicator that snooping is disabled is that we see the packets in the control-plane capture, but not in the DHCP snooping debug.

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 3: DHCP snooping is disabled on the Edge node.

What happens when DHCP snooping is disabled? The Agent Remote ID sub-option of Option 82 is not inserted.

```
Edge1#show monitor capture cap buffer display-filter frame.number==53 detail
…
Dynamic Host Configuration Protocol (Offer)
    Message type: Boot Reply (2)
…
Option: (82) Agent Information Option
        Length: 8
        Option 82 Suboption: (1) Agent Circuit ID
            Length: 6
            Agent Circuit ID: 000403fe0103
…
```

There is no Agent Remote ID, which is the sub-option of Option 82 that contains the IP of the Loopback 0 of the Edge that originated the DISCOVER.

```
Edge1#show ip dhcp snooping
Switch DHCP snooping is disabled
Switch DHCP gleaning is disabled
DHCP snooping is configured on following VLANs:
11,101,1021,1025-1026 <<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<< The VLAN 1022 is not present.
DHCP snooping is operational on following VLANs:
11,101,1021,1025-1026…
```

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 3: DHCP snooping is disabled on the Edge node.

Viewing the OFFER packet on the Border, the Agent Remote ID sub-option is also not present. Which Border will the Fusion send the OFFER to?

```
Fusion#show ip cef 172.16.20.1
172.16.20.0/24
  nexthop 172.16.254.9 Vlan82
```

The OFFER will be sent to the next hop 172.16.254.9; this is the SVI used for eBGP on Border1.

```
Border1#monitor capture cap control-plane both match any buffer size 100 start
Border1#sh monitor capture cap buffer display-filter "bootp and ip.addr==172.16.20.1"
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

  395   7.577153 172.16.0.1 -> 172.16.20.1  DHCP 373 DHCP Offer    - Transaction ID 0x26dc
  480  10.881252 172.16.0.1 -> 172.16.20.1  DHCP 373 DHCP Offer    - Transaction ID 0x26dc
```

The Agent Remote ID sub-option is missing.

```
Border1#show monitor capture cap buffer display-filter frame.number==480 detail
…
Dynamic Host Configuration Protocol (Offer)
    Message type: Boot Reply (2)
…
Option: (82) Agent Information Option
        Length: 8
        Option 82 Suboption: (1) Agent Circuit ID
            Length: 6
            Agent Circuit ID: 000403fe0103
```

© 2024 Cisco and/or its affiliates. All rights reserved.

Problem 3: DHCP snooping is disabled on the Edge node.

How should Option 82 look in an SDA environment?

```
Edge1#show monitor capture cap buffer display-filter frame.number==33 detail
…
Option: (82) Agent Information Option
        Length: 20
        Option 82 Suboption: (1) Agent Circuit ID
            Length: 6
            Agent Circuit ID: 000403fe0103
        Option 82 Suboption: (2) Agent Remote ID
            Length: 10
            Agent Remote ID: 030800100601ac100145 <<<<< RLOC is the last 8 hexadecimal values
```

Loopback0 on Edge1

```
Edge1#show run int loopback 0
interface Loopback0
 description Fabric Node Router ID
 ip address 172.16.1.69 255.255.255.255
```

If we convert the hexadecimal value 01ac100145 to an IP address, we get the RLOC IP of the Edge, which is 172.16.1.69.

© 2024 Cisco and/or its affiliates. All rights reserved.

Q&A

© 2024 Cisco and/or its affiliates. All rights reserved.

Our social Media

LinkedIn
[Cisco Community]

Twitter
[@CiscoCommunity]

YouTube
[CiscoCommunity]

Facebook
[CiscoCommunity]

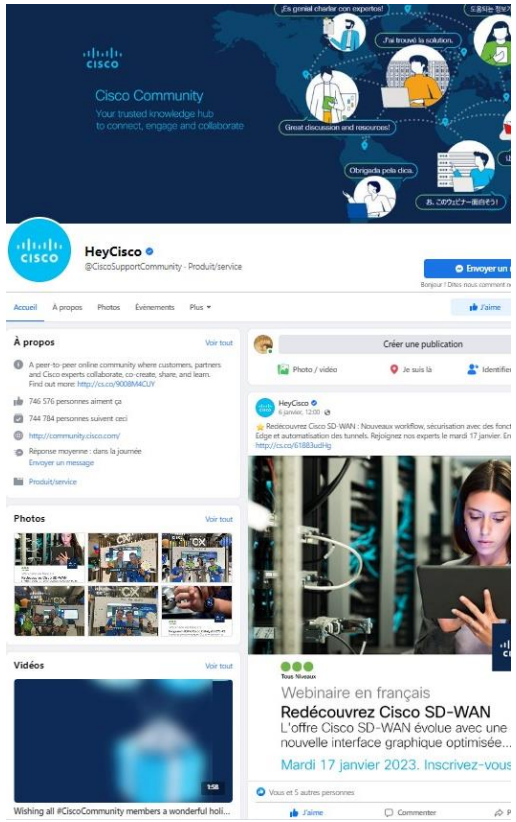

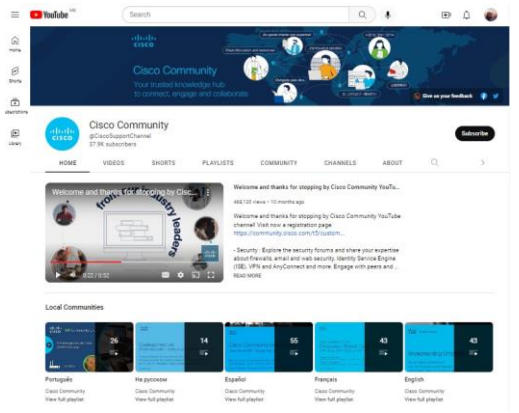

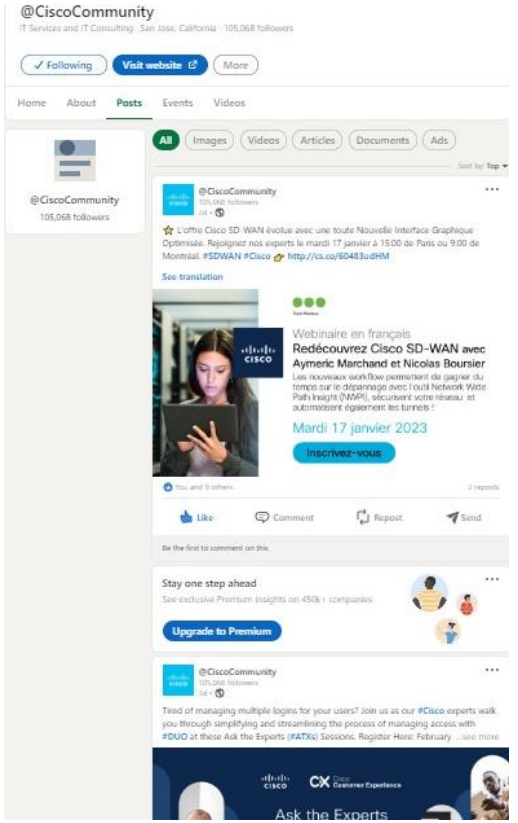

© 2024 Cisco and/or its affiliates. All rights reserved.


CISCO