
cisco
You make possible

Cisco Software-Defined Access

Technology Deep Dive

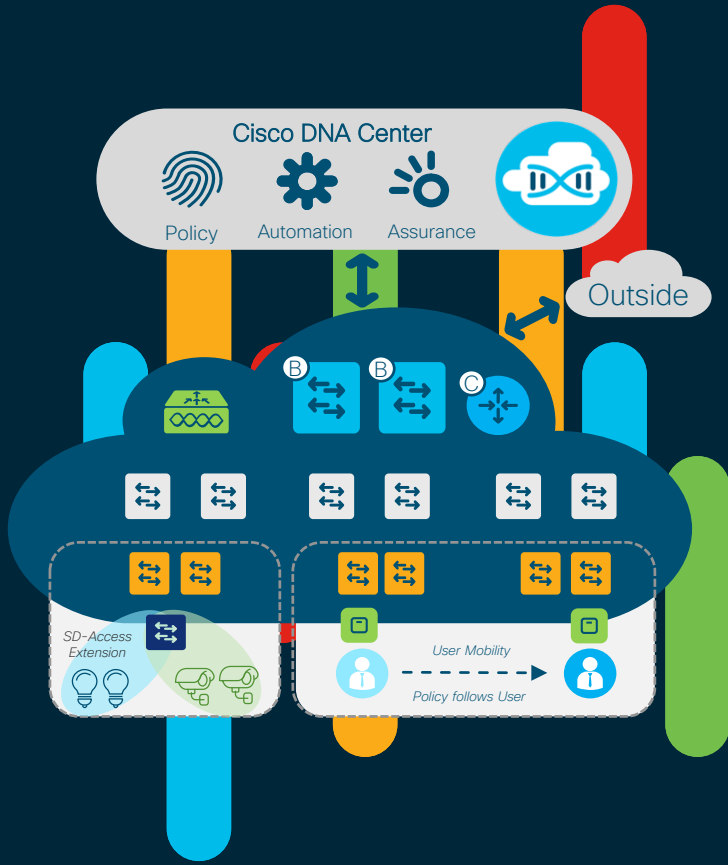

Victor Moreno
Distinguished Engineer - Technical Marketing

BRKCRS-3810

cisco Live!

Barcelona | January 27-31, 2020

Cisco Webex Teams

Questions?

Use Cisco Webex Teams to chat with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

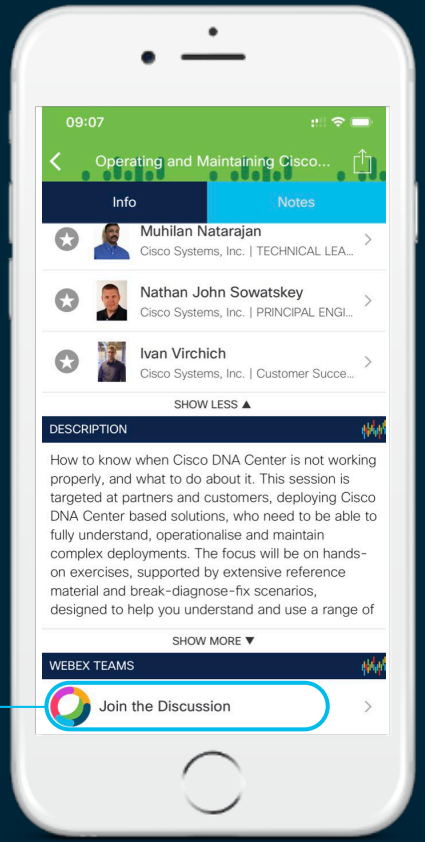

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Assumptions

This session assumes you have received Cisco DNA, SD-Access & ISE Training

If not… please complete one or all of the following training materials:

- CiscoLive
- dCloud Lab
- Learning@Cisco
- SDA Design CVD
- SDA Deploy CVD
- DNAC Guides

This session is based on Cisco DNAC / SDA 1.3.1, ISE 2.6 and IOS-XE 16.11

- Product Compatibility Matrix

For a full list of current capabilities, restrictions, limitations & caveats refer to:

- DNAC Release Notes

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- Cisco SD-Access Recap
- Host Onboarding
  - Endpoint Classification
  - DHCP in SD-Access Fabric
- Connectivity & Access Control
  - Unicast Forwarding
  - Access Control Policy
- Advanced Topics
  - Multicast Forwarding
  - Broadcast Forwarding

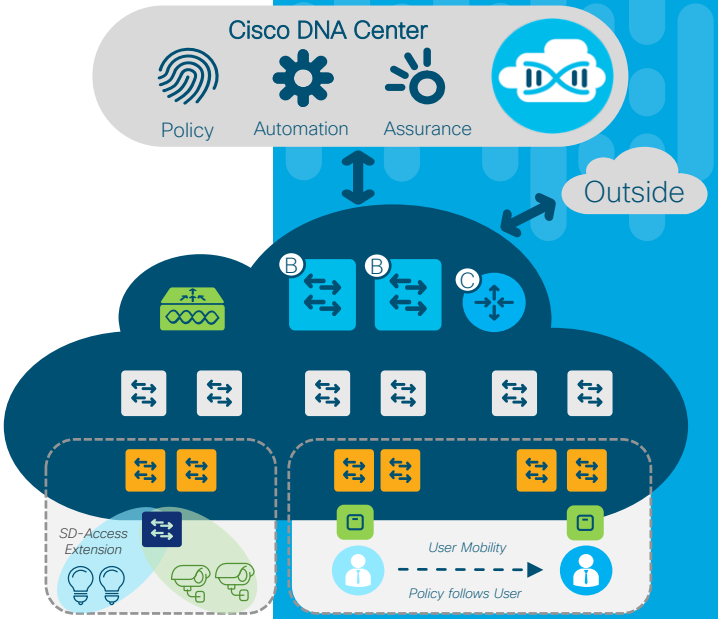

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

1. Architecture
2. Key Components
3. Fabric Constructs

What is Campus Fabric?

Fabric Fundamentals

Cisco Software Defined Access

The Foundation for Cisco's Intent-Based Network

RECAP

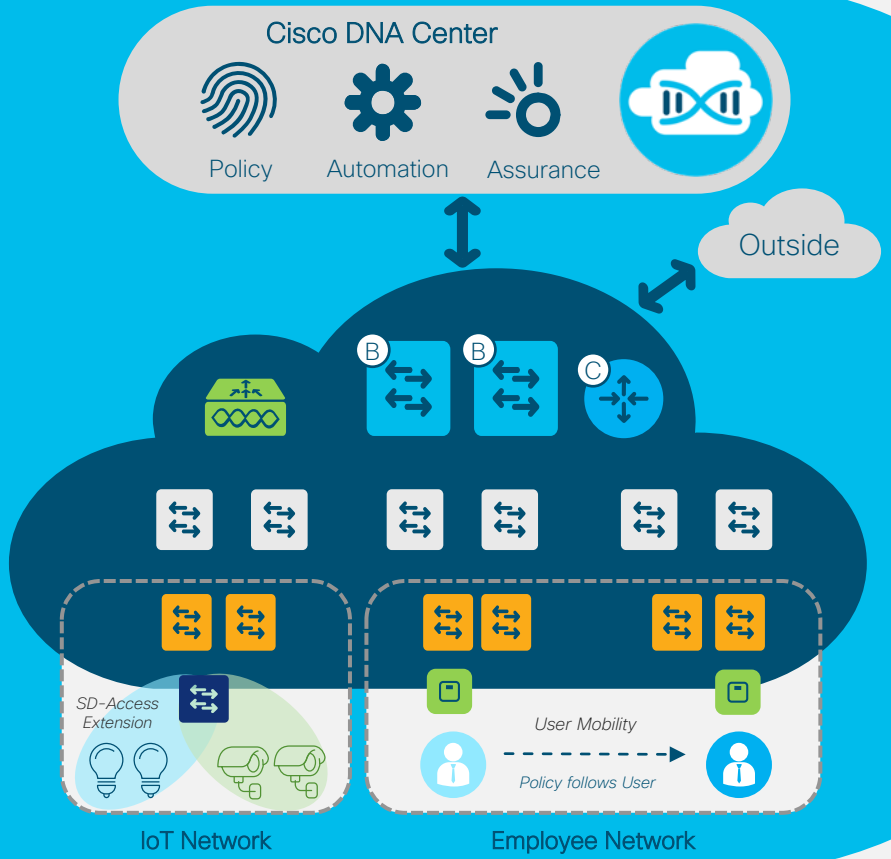

Identity-Based Policy and Segmentation

Policy definition decoupled from VLAN and IP address

Automated Network Fabric

Single fabric for Wired and Wireless with full automation

Insights and Telemetry

Analytics and insights into User and Application experience

SD-Access Architecture

Fabric Roles & Terminology

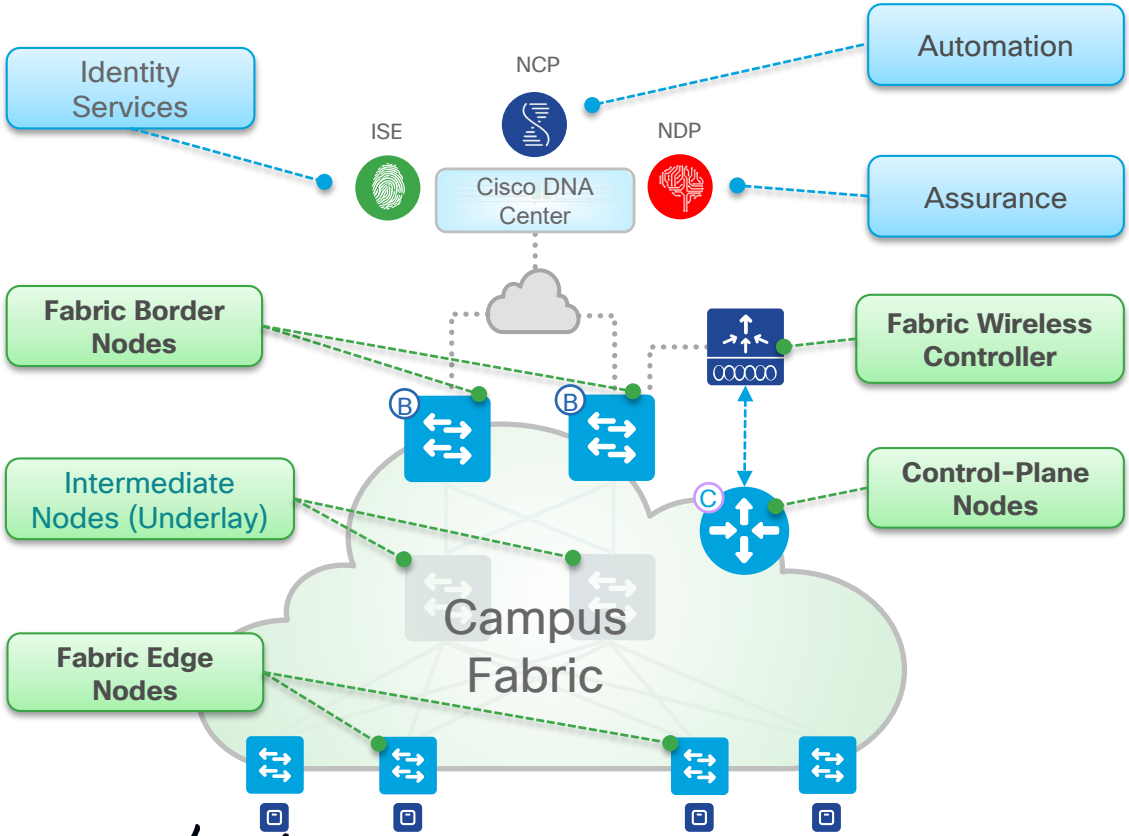

RECAP

- **Network Automation** – Simple graphical user interface and intent based automation (e.g. NCP) of fabric devices
- **Network Assurance** – Data Collectors (e.g. NDP) analyze Endpoint to App flows and monitor fabric status
- **Identity Services** – NAC & ID Systems (e.g. ISE) for dynamic Endpoint to Group mapping and Policy definition
- **Control-Plane Nodes** – Map System that manages Endpoint to Device relationships
- **Fabric Border Nodes** – A Fabric device (e.g. Core) that connects External L3 network(s) to the SD-Access Fabric
- **Fabric Edge Nodes** – A Fabric device (e.g. Access or Distribution) that connects Wired Endpoints to the SD-Access Fabric
- **Fabric Wireless Controller** – A Fabric device (WLC) that connects APs and Wireless Endpoints to the SD-Access Fabric

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

Campus Fabric - Key Components

RECAP

1. **Control-Plane** based on **LISP**
2. **Data-Plane** based on **VXLAN**
3. **Policy-Plane** based on **CTS**

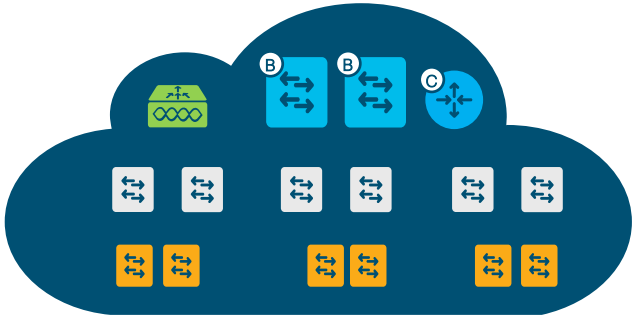

Key Differences

- L2 + L3 Overlay -vs- L2 or L3 Only
- Host Mobility with Anycast Gateway
- Adds VRF + SGT into Data-Plane
- Virtual Tunnel Endpoints (Automatic)
- NO Topology Limitations (Basic IP)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Locator / ID Separation Protocol

Location and Identity separation

RECAP

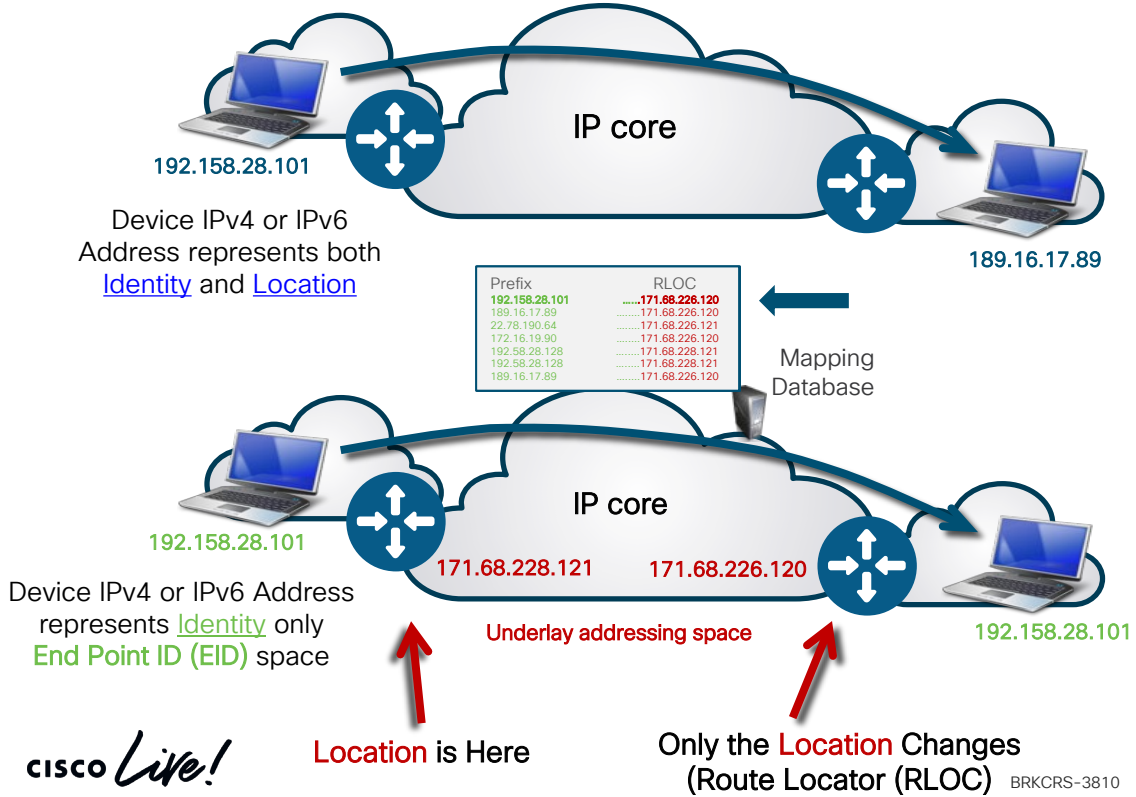

Traditional Behavior – Location + ID are "Combined"

When the Device moves, it gets a new IPv4 or IPv6 Address for its new Identity and Location

Overlay Behavior – Location & ID are "Separated"

When the Device moves, it keeps the same IPv4 or IPv6 Address. It has the Same Identity

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

LISP Control Plane

RECAP

Fabric nodes use LISP as a control plane for Endpoint Identifier (EID) and Routing Locator (RLOC) info

Fabric Control Plane node acts as a Map Server / Resolver for EID to RLOC mappings

Fabric Edge and Internal Border devices registers EIDs to the Map Server.

External Border node acts as PXTR (LISP Proxy Tunnel Router) and provides default gateway when no mapping exists.

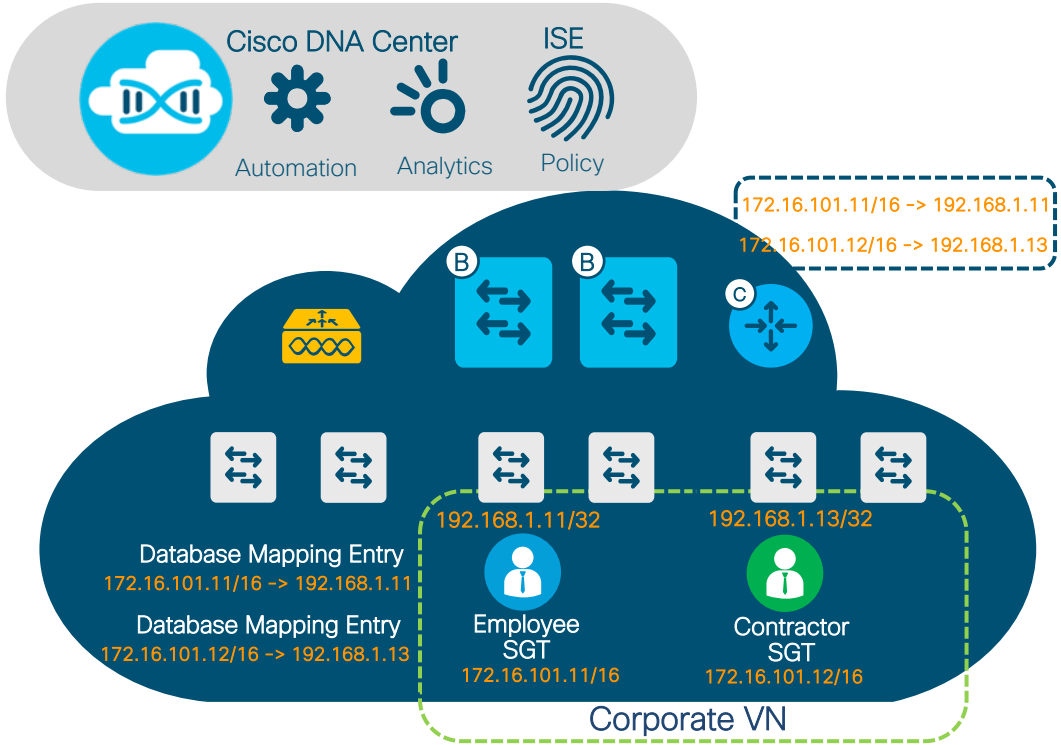

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

VXLAN Data Plane

RECAP

Fabric nodes use VXLAN (Ethernet Based) as the data plane which supports both L2 and L3 overlay.

VXLAN header contains VNID (VXLAN Network Identifier) field which allows up to 16 million Virtual Networks (Layer 3 (VRFs) or Layer 2).

VXLAN header also has Group Policy ID, or Scalable Group Tags (SGTs), allowing 64,000 SGTs.

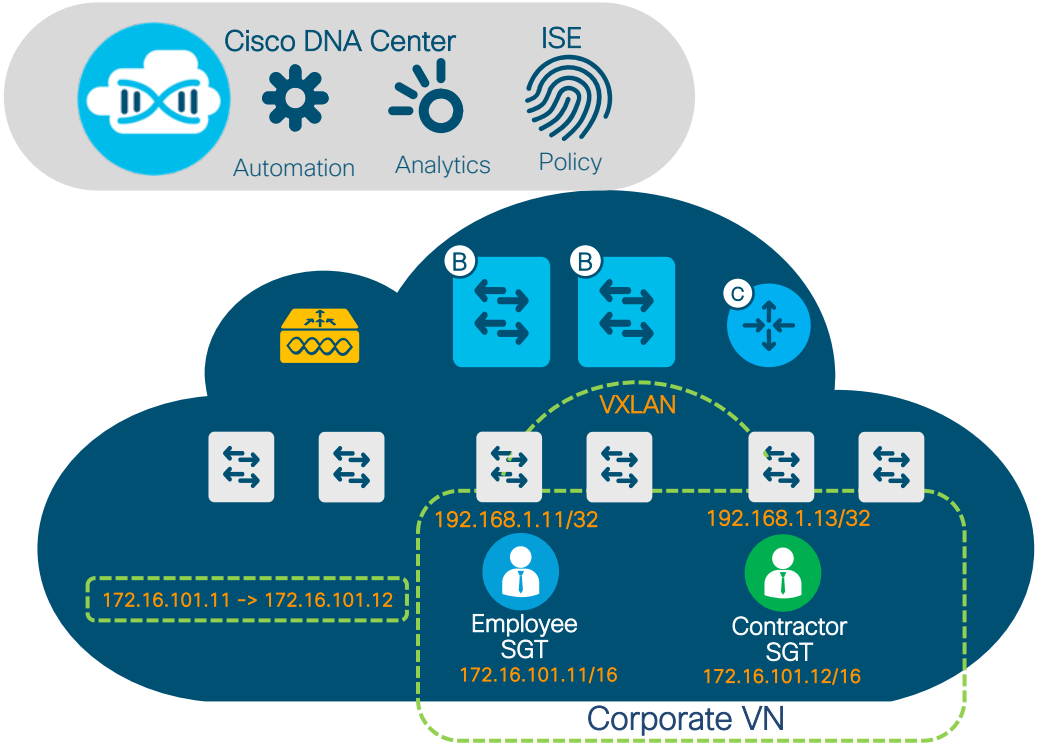

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

Cisco TrustSec Policy Plane

RECAP

Scalable Groups (SG) are a logical construct defined/identified based on the user and/or device context.

ISE dynamically assign SGs to the users and devices coming to the network fabric.

Nodes add Scalable Group Tags (SGTs) to the fabric encapsulation when communicating between the users and devices.

Edge and border nodes enforce the SGACL policies for the SGTs they protect locally.

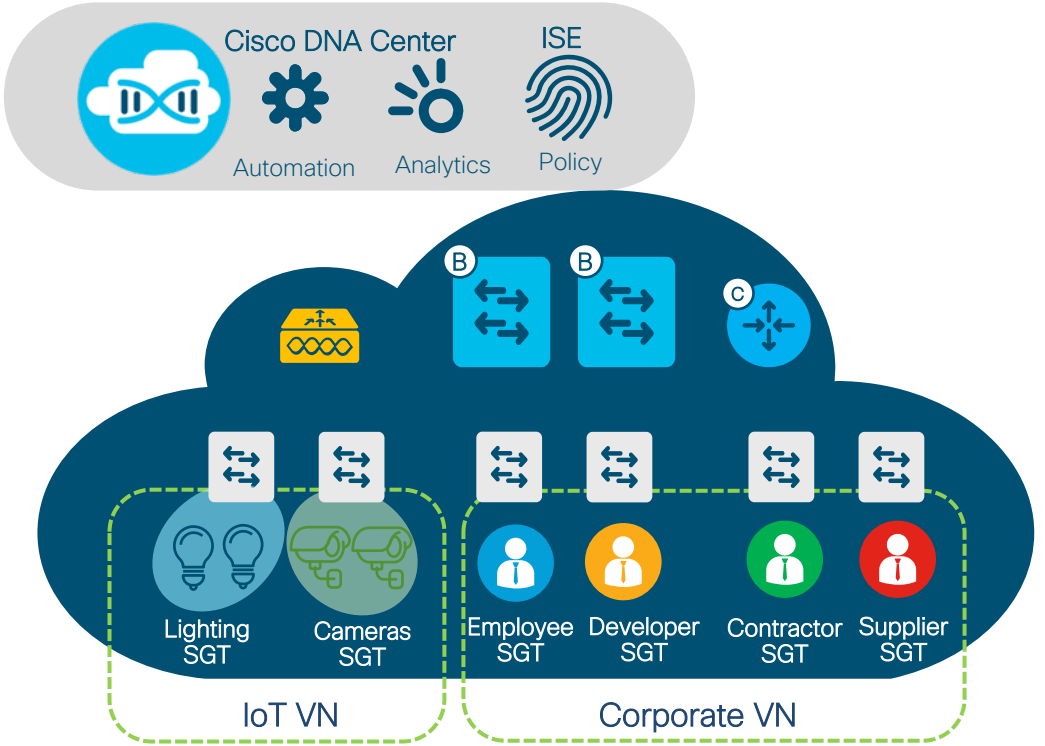

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

How VNs work in SD-Access

RECAP

- **Fabric Devices (Underlay)** connectivity is in the Global Routing Table
- **INFRA_VN** is only for Access Points and Extended Nodes in GRT
- **DEFAULT_VN** is an actual "User VN" provided by default
- **User-Defined VNs** can be added or removed on-demand

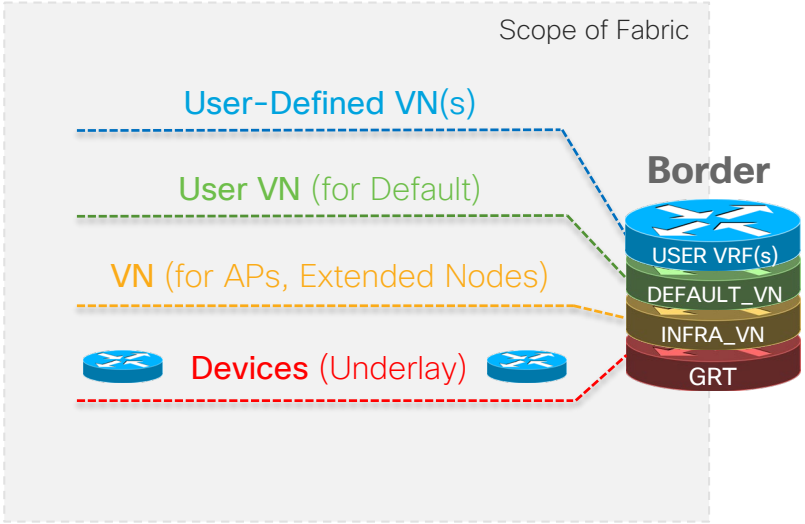

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric

How VNs work in SD-Access

RECAP

SD-Access Designs connecting to existing Global Routing Table should use a “Fusion” router with MP-BGP & VRF import/export.

```
ip vrf USERS
 rd 1:4099
 route-target export 1:4099
 route-target import 1:4099
 route-target import 1:4097
!
ip vrf DEFAULT_VN
 rd 1:4098
 route-target export 1:4098
 route-target import 1:4098
 route-target import 1:4097

ip vrf GLOBAL
 rd 1:4097
 route-target export 1:4097
 route-target import 1:4097
 route-target export 1:4099
 route-target export 1:4098
```

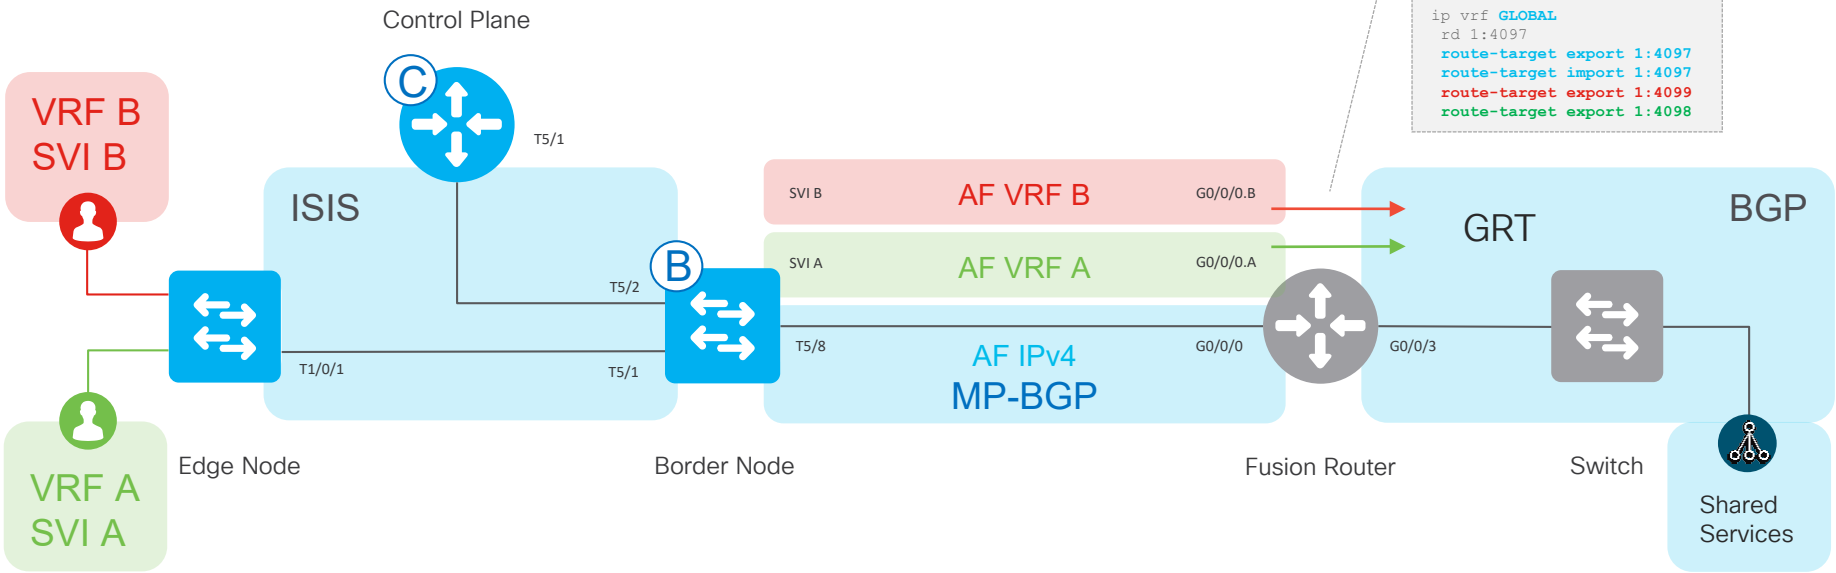

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

ISE – Cisco DNA Center Operation

RECAP

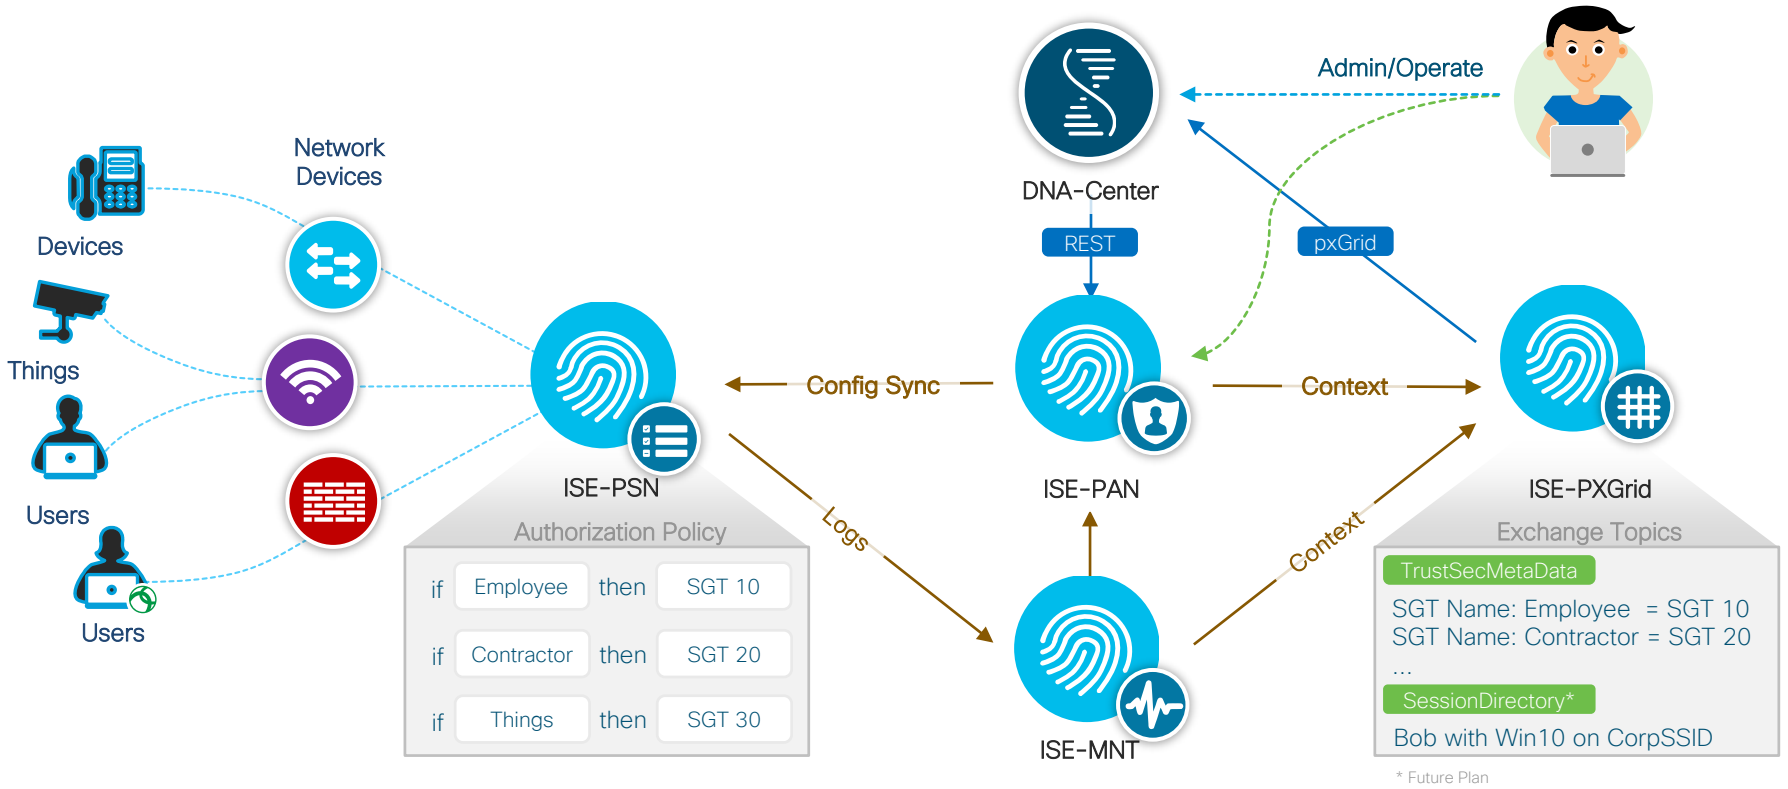

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Policy

Two Level Hierarchy - Macro Level

RECAP

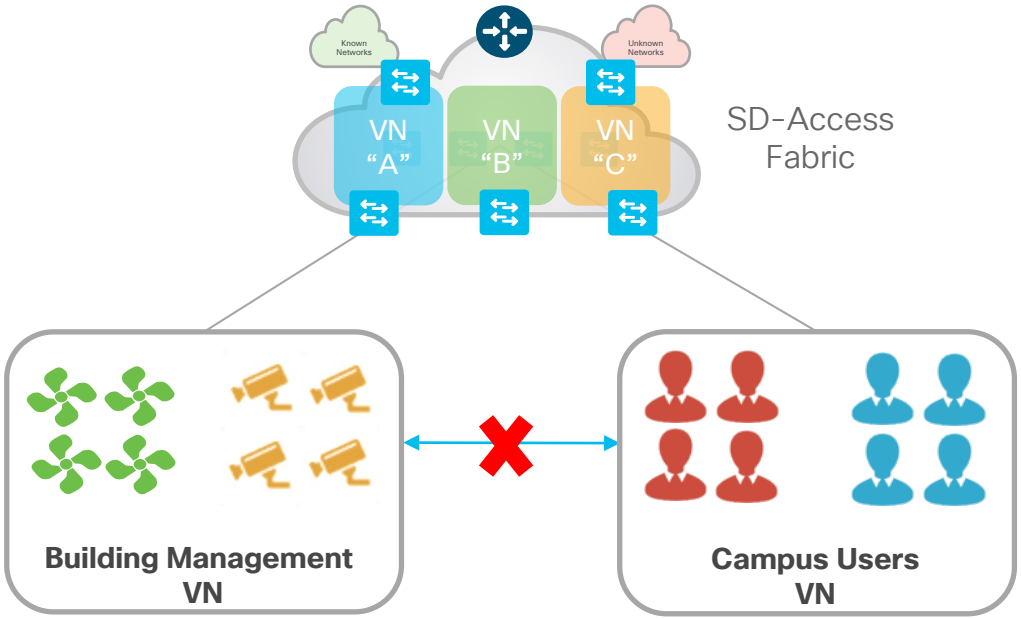

Virtual Network (VN)

First level Segmentation ensures **zero communication** between forwarding domains. Ability to consolidate multiple networks into one management plane.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Policy

Two Level Hierarchy - Micro Level

RECAP

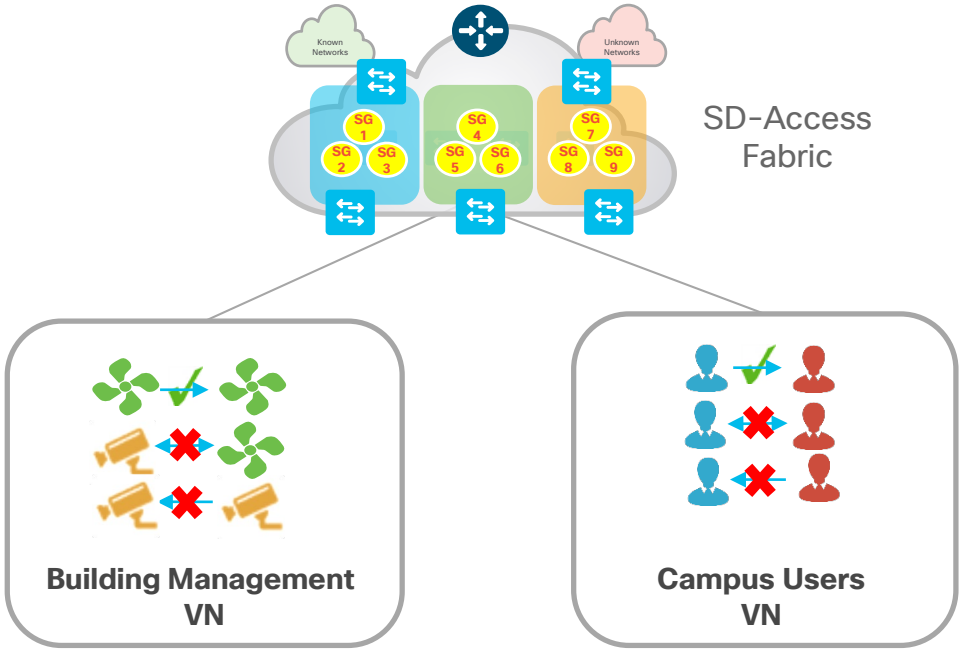

Scalable Group (SG)

Second level Segmentation ensures **role based access control** between two groups within a Virtual Network. Provides the ability to segment the network into either line of businesses or functional blocks.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

1. Endpoint Classification
2. DHCP in Fabric

Classification and Group tagging

Host Onboarding

Authentication and Authorization

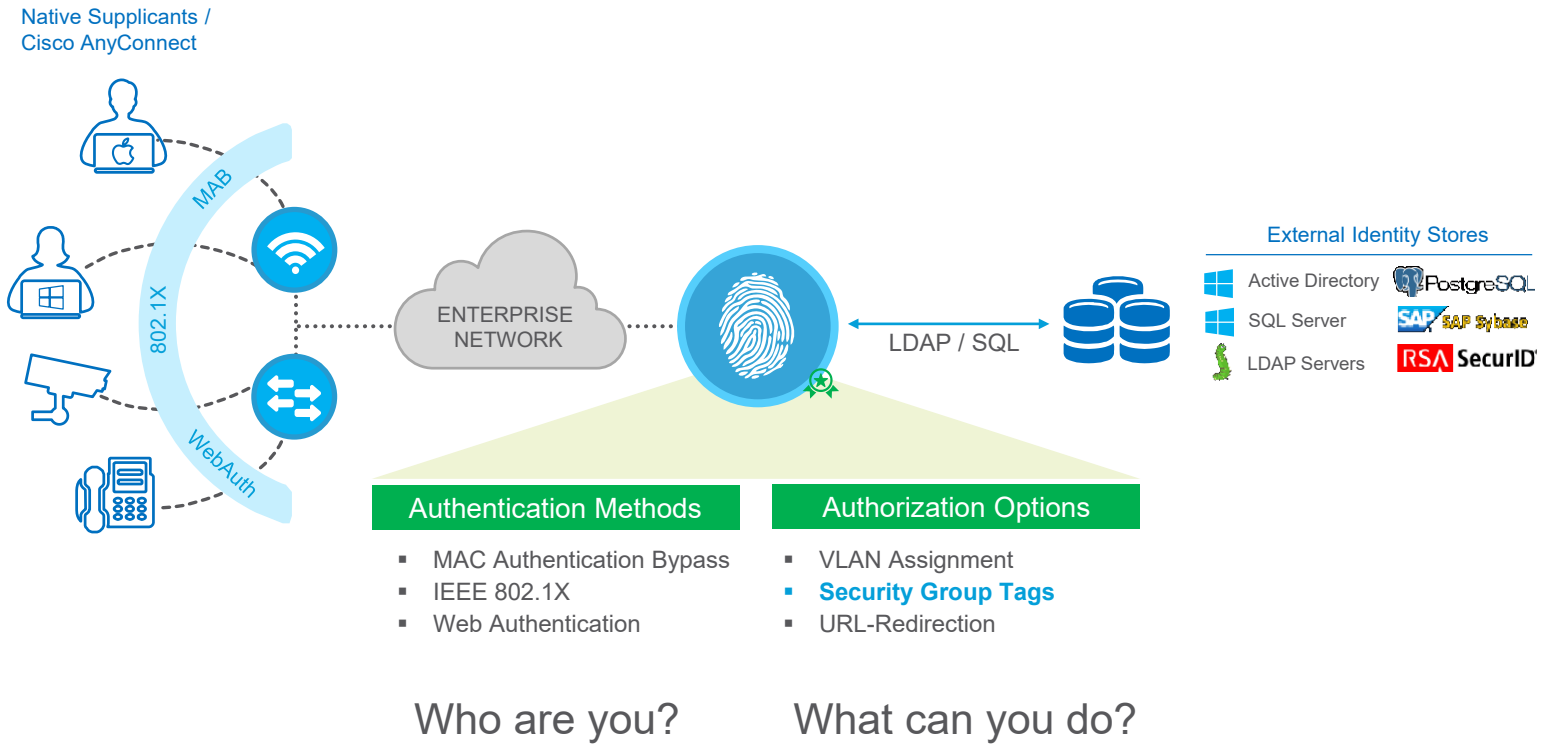

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Authorization Workflow for Segmentation

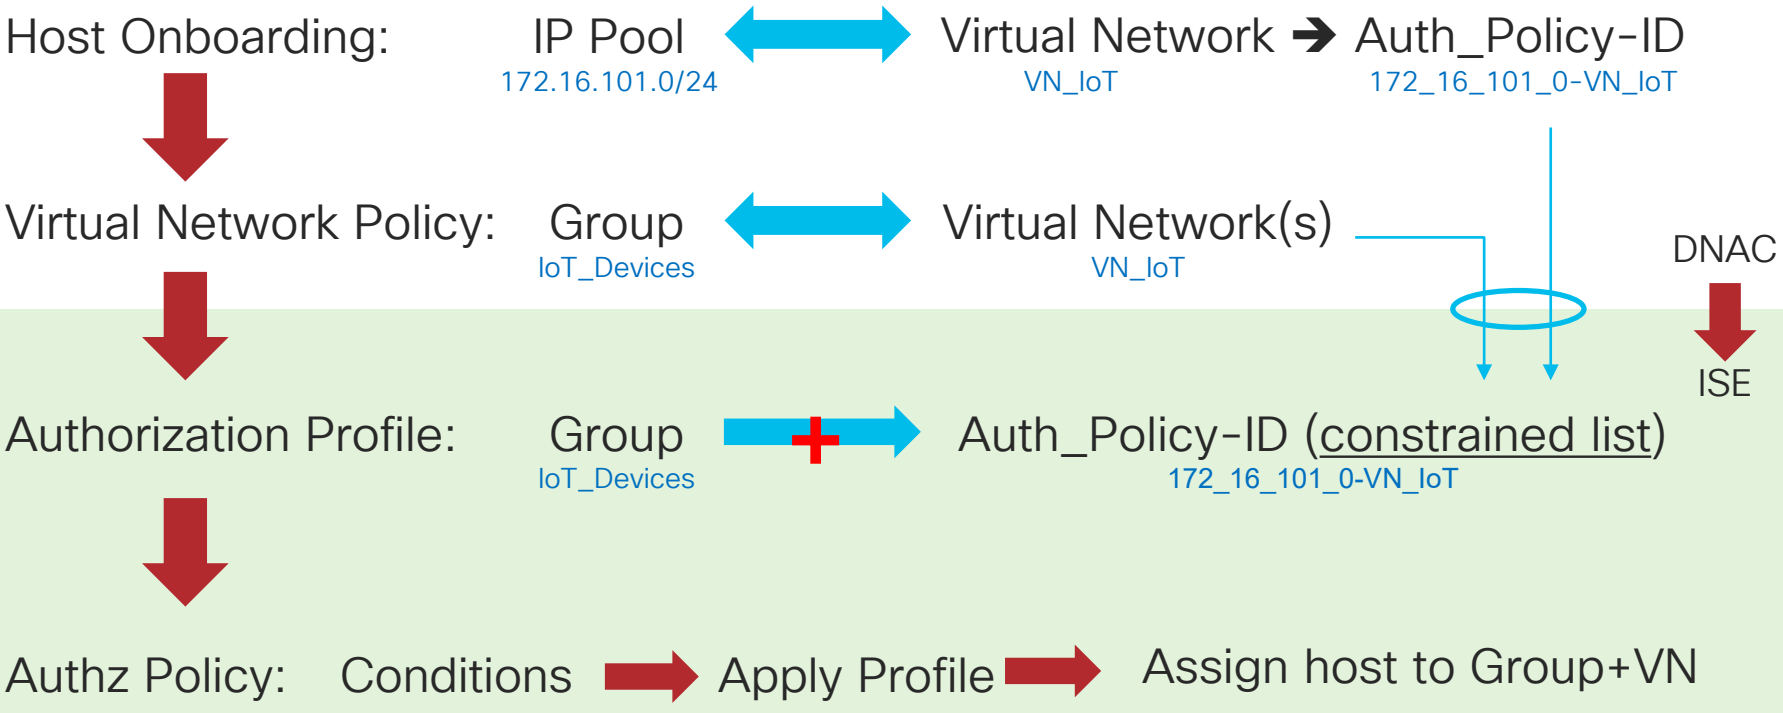

Authorization Policy – Authorization Profiles

Policy Sets | Profiling | Posture | Client Provisioning | Policy Elements

Dictionaries | Conditions | Results

- Authentication
- Authorization
  - Authorization Profiles
  - Downloadable ACLs
- Profiling
- Posture
- Client Provisioning

Authorization Profiles > **VN_IoT_w_SGT**

Authorization Profile

* Name: VN_IoT_w_SGT

Description:

* Access Type: ACCESS_ACCEPT

Network Device Profile: Cisco

Service Template

Track Movement

Passive Identity Tracking

Common Tasks

DACL Name

ACL (Filter-ID)

Security Group: IOT_Devices — Virtual Network: VN_IoT

VLAN

Advanced Attributes Settings

Select an item =

Attributes Details

```
Access Type = ACCESS_ACCEPT
cisco-av-pair = cts:security-group-tag=0010-0
cisco-av-pair = cts:sgt-name=IOT_Devices
cisco-av-pair = cts:vn= VN_IoT
Tunnel-Private-Group-ID = 1:
Tunnel-Type = 1:13
Tunnel-Medium-Type = 1:6
```


© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

ISE Authorization for VN assignment

Authorization Result = Virtual Network + Security Group

Work Centers > TrustSec > Authorization Policy

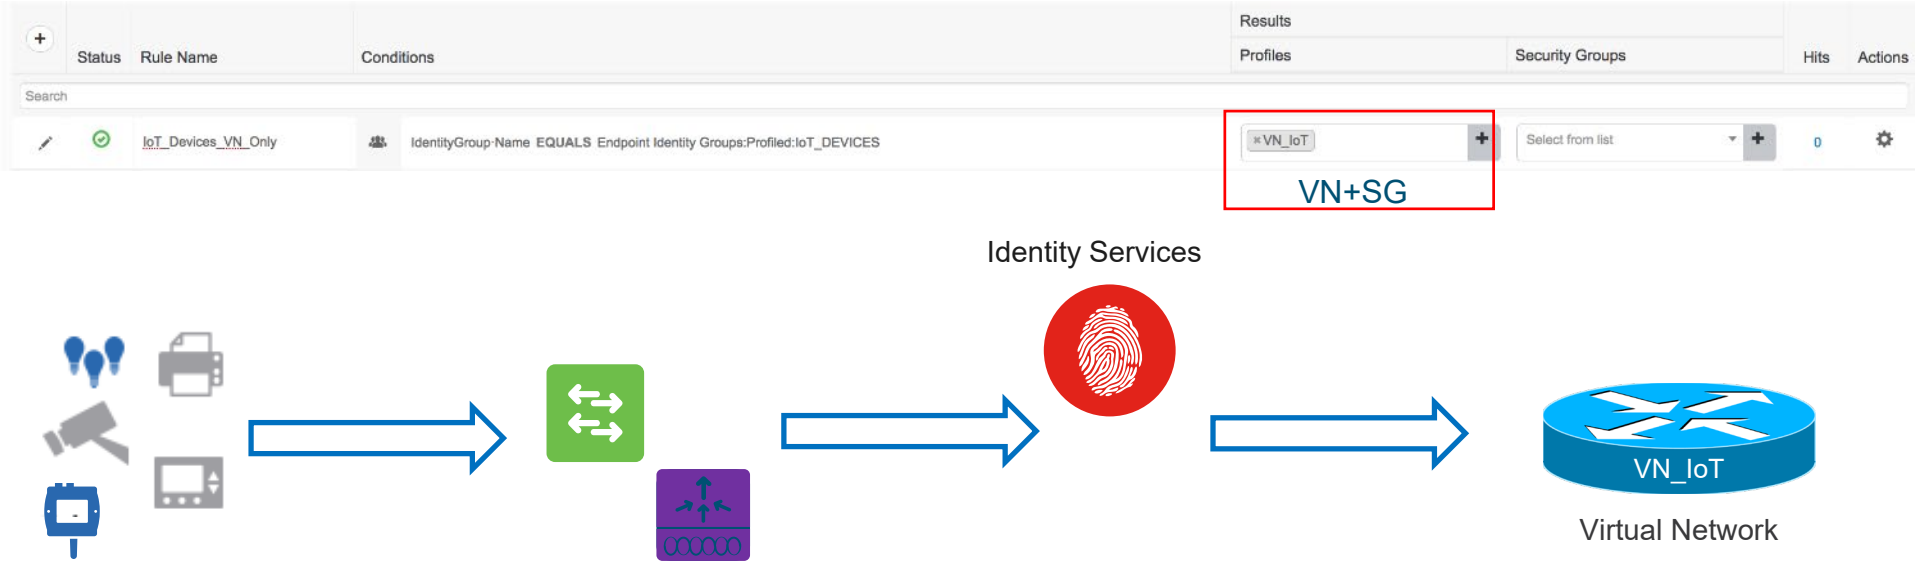

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Group Assignment

Two ways to assign SGT

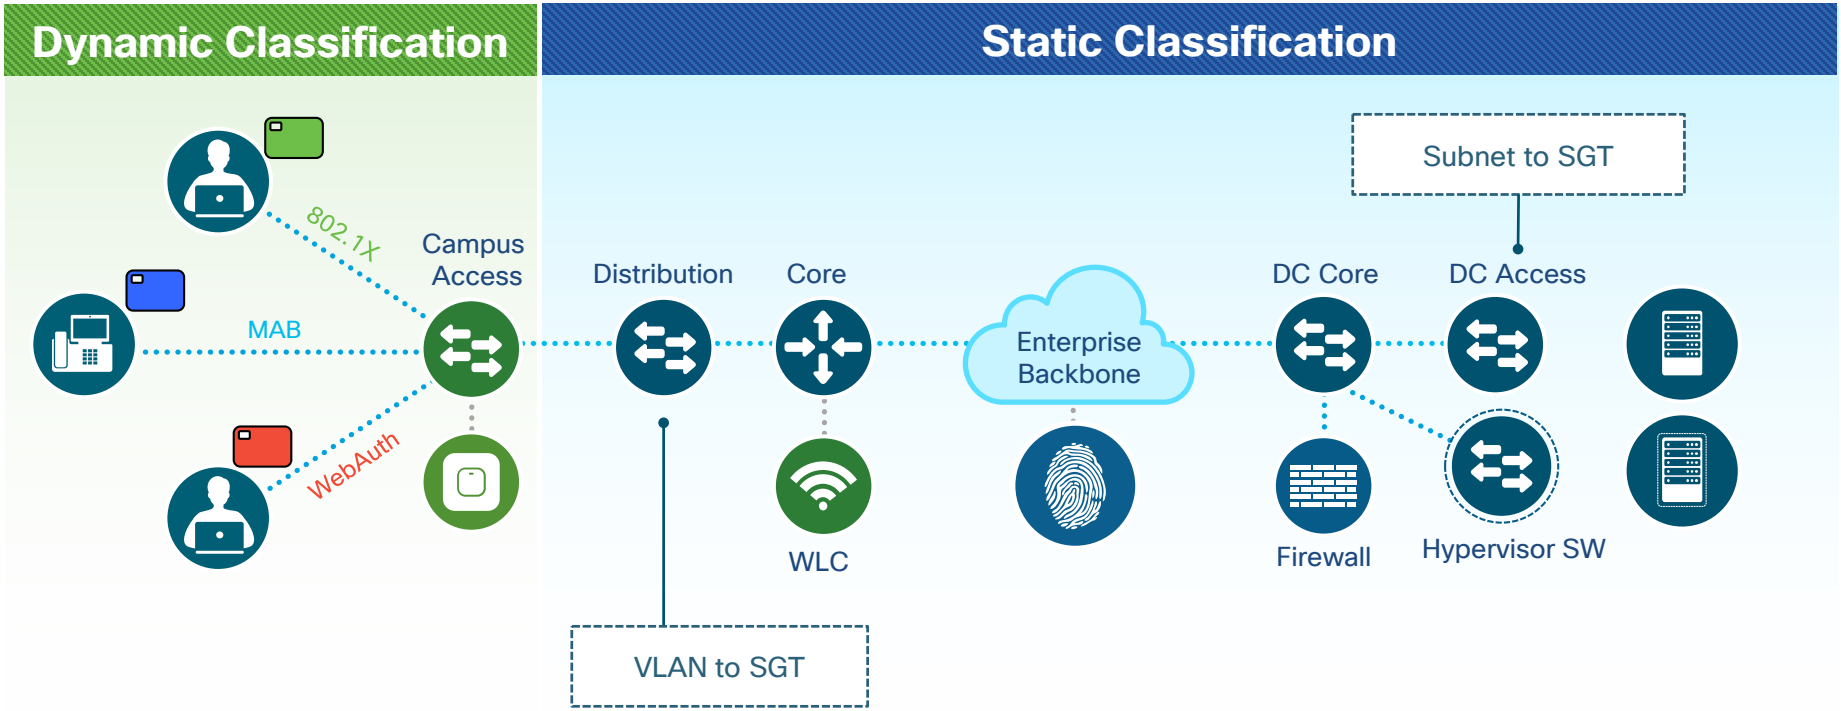

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Defining Security Groups

Define SGs under '**Components**' section in **TrustSec Work Center** (from ISE 2.0)

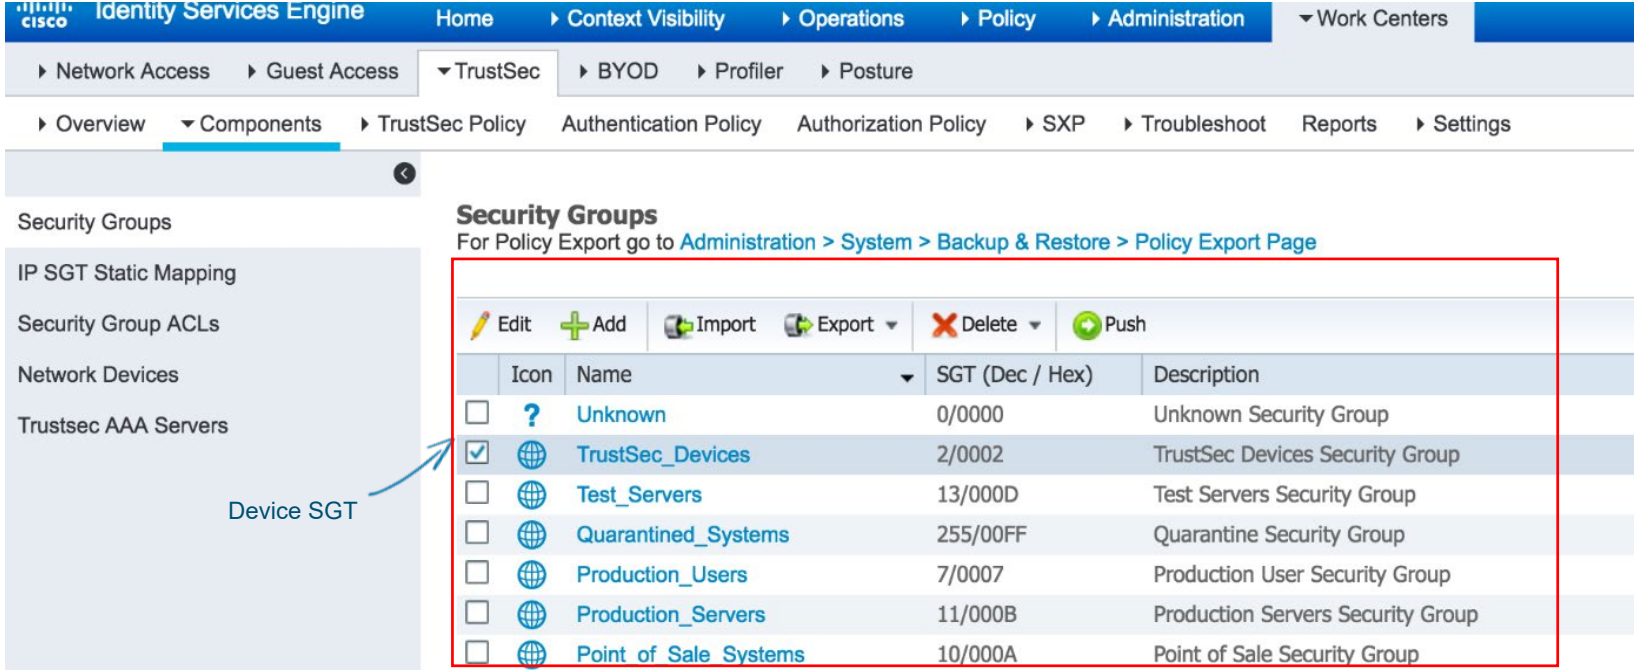

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SGT for wired 802.1X session

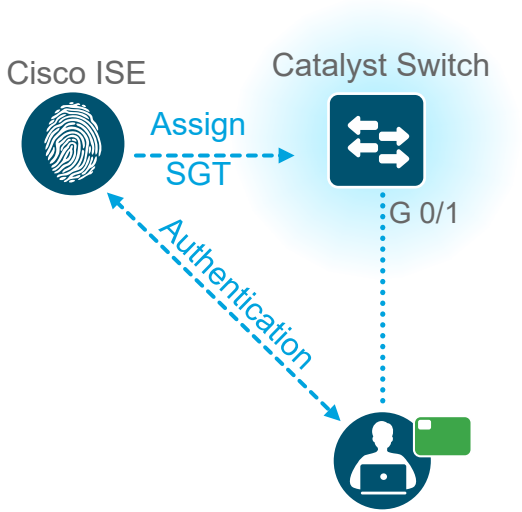


```
Switch# show authentication sessions int Gi 0/1 details
            Interface:  GigabitEthernet1/0/23
              IIF-ID:  0x107AB4000000076
         MAC Address:  0005.0005.0005
        IPv6 Address:  2001:DB8:100:0:3809:A879:5197:16DB
        IPv4 Address:  172.20.100.2
           User-Name:  bob@trustsec.lab
              Status:  Authorized
              Domain:  DATA
      Oper host mode:  multi-auth
    Oper control dir:  both
     Session timeout:  N/A
   Common Session ID:  0A01010100000FC50BEC5800
     Acct Session ID:  0x00000FBE
              Handle:  0xD4000009
      Current Policy:  POLICY_Gi1/0/23

Server Policies:
            SGT Value:  10

Method status list:
       Method           State
       mab              Authc Success
```

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SGTs for Wireless sessions

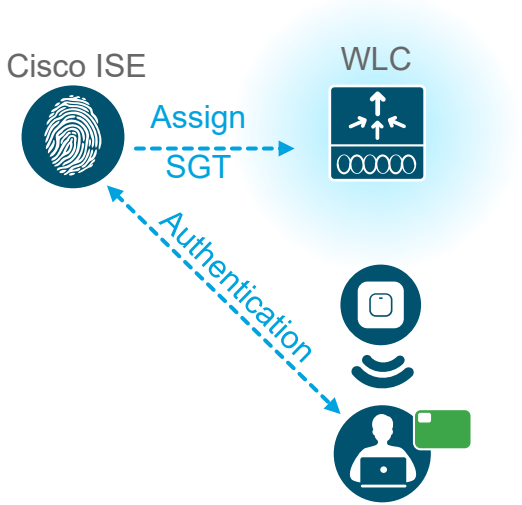

Works on AirOS and IOS-XE Wireless controllers.

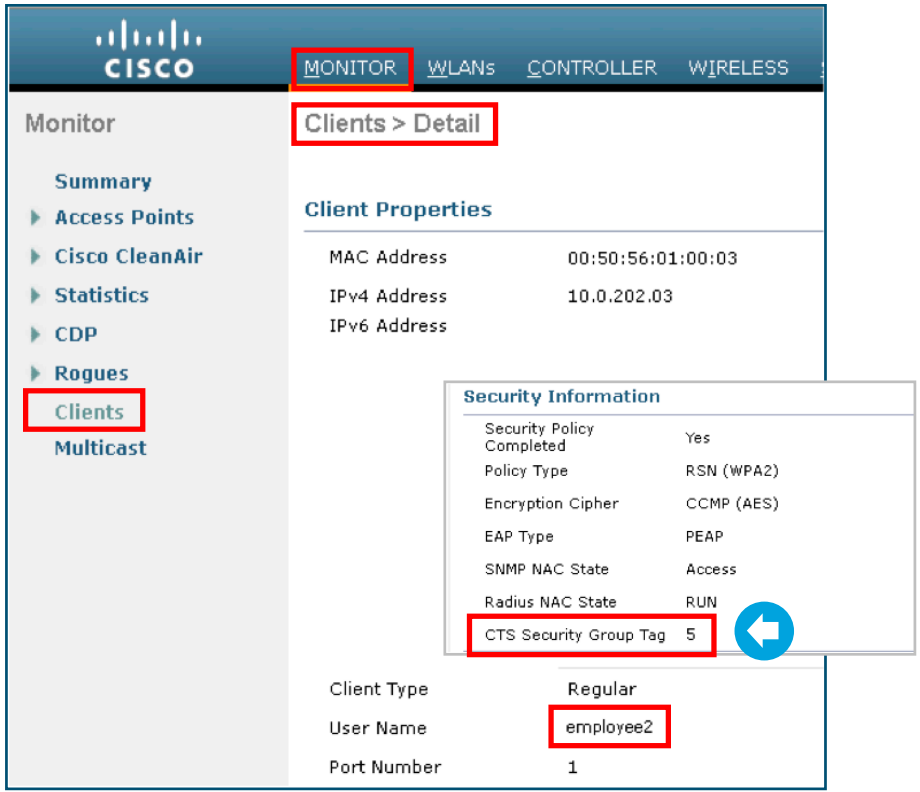

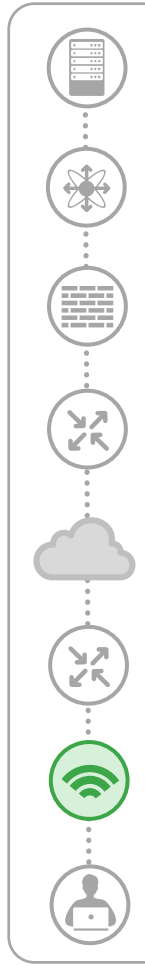

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Gaining access to network fabric

Host Onboarding

1. Endpoint Classification
2. DHCP in Fabric

SD-Access Fabric Architecture

DHCP reachability

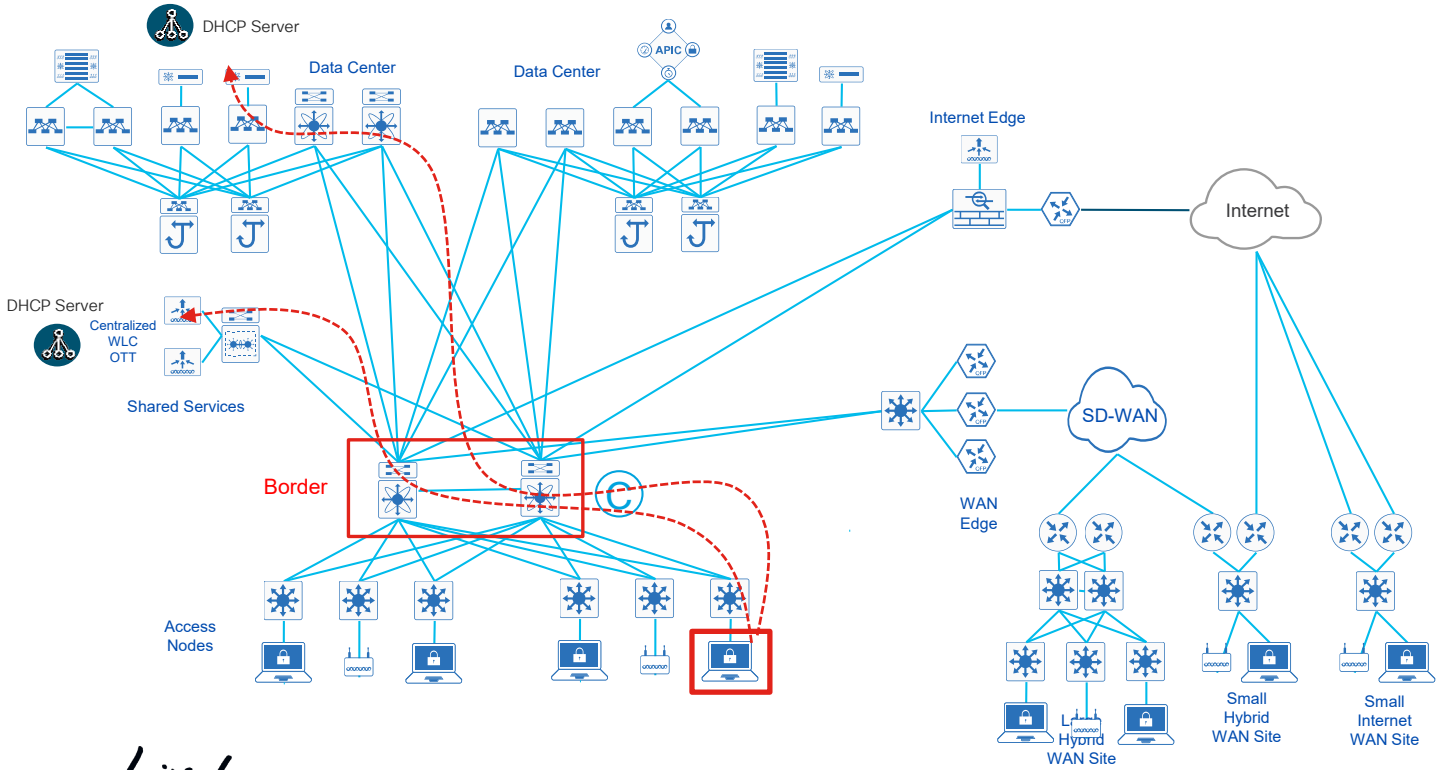

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in an anycast Gateway environment

RECAP

Anycast GW provides a single L3 Default Gateway for IP capable endpoints

- The same Switch Virtual Interface (SVI) is present on EVERY Edge with the SAME Virtual IP and MAC
- When a Host moves from Edge 1 to Edge 2, it does not need to change its Default Gateway ☺
- The SVI is also configured with an IP helper address for DHCP.

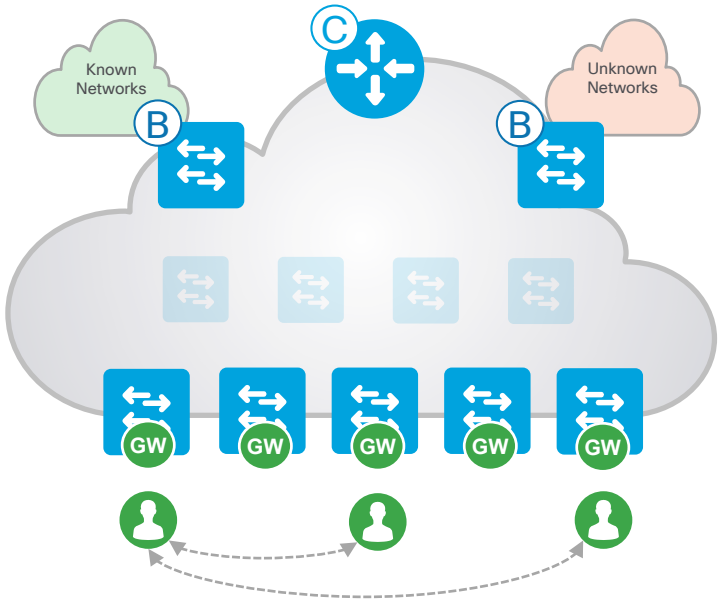

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in an anycast Gateway environment

DHCP reply needs to come to the right edge node switch.

- But “we do not know on which edge node a host is located” as we don’t have an IP address for it yet.
- Once an IP address is assigned to the host, the control plane node learns where the host is located.
- The Control-Plane maintains the Host to Edge relationship (Fabric Dynamic EID mapping)

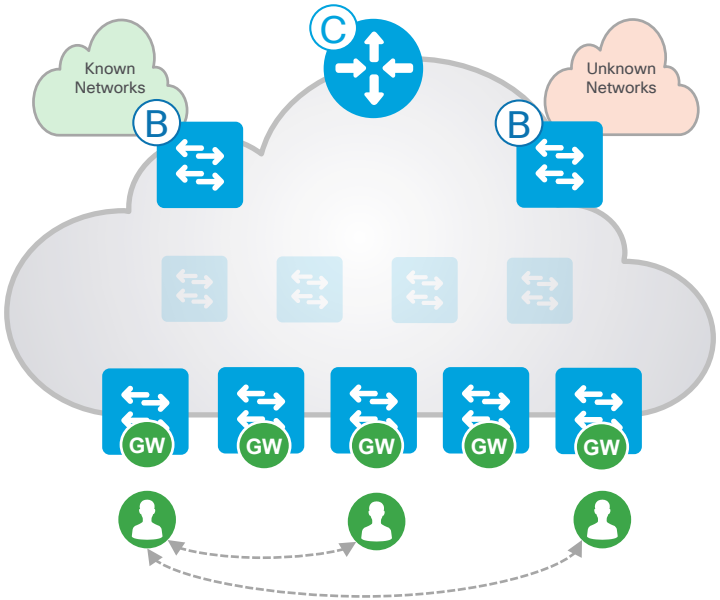

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – Enable LISP lookups for DHCP requests

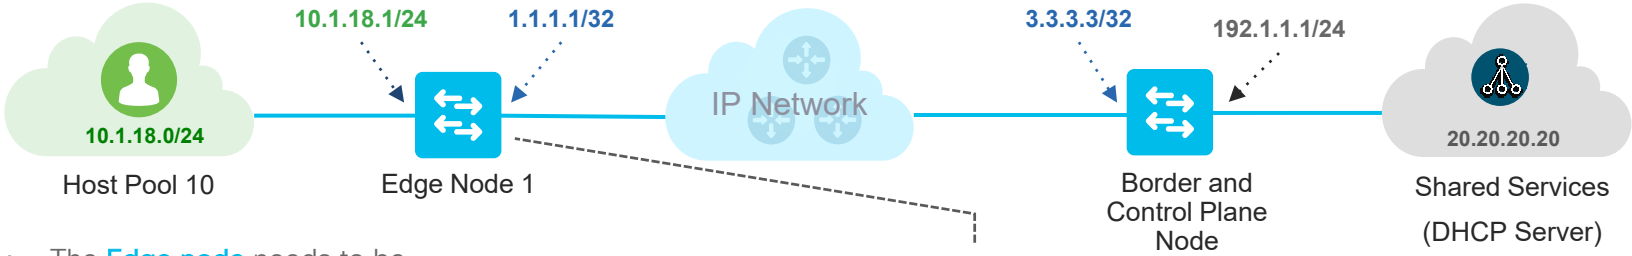

- The Edge node needs to be configured as a "proxy-itr" to avoid source EID validation. The source of the DHCP request doesn't have an IP address yet.
- We also need a 0/0 map-cache that triggers a LISP lookup for the DHCP helper address so that the DHCP request is sent in the overlay.

```
router lisp
instance-id 4098
  dynamic-eid user
   database-mapping 10.1.18.0/24 locator-set edge1
   exit-dynamic-eid
  !
  service ipv4
   eid-table vrf User
   map-cache 0.0.0.0/0 map-request
   itr map-resolver 3.3.3.3
   proxy-itr 1.1.1.1
   etr map-server 3.3.3.3 key uci
   etr
   use-petr 3.3.3.3
   exit-service-ipv4
  !
  exit-instance-id
!
exit-router-lisp
```


© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – Option 82 plus Snooping

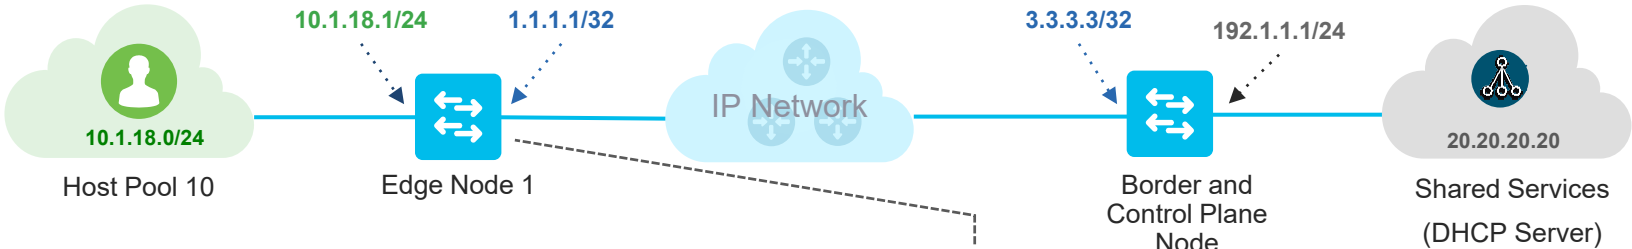

- The Edge node needs to enable option 82 in the DHCP request.
  - Option 82 will carry the VNID and RLOC.
- DHCP snooping needs to be enabled on all the VLANs in fabric.

```
ip dhcp relay information option
ip dhcp snooping vlan 1022
ip dhcp snooping

interface Vlan1022
ip vrf forwarding User
ip address 10.1.18.1 255.255.255.0
ip helper-address 20.20.20.20
lisp mobility user
```

SD-Access Fabric Architecture

DHCP in Fabric – "Pre-advertise" IP pools outside the fabric

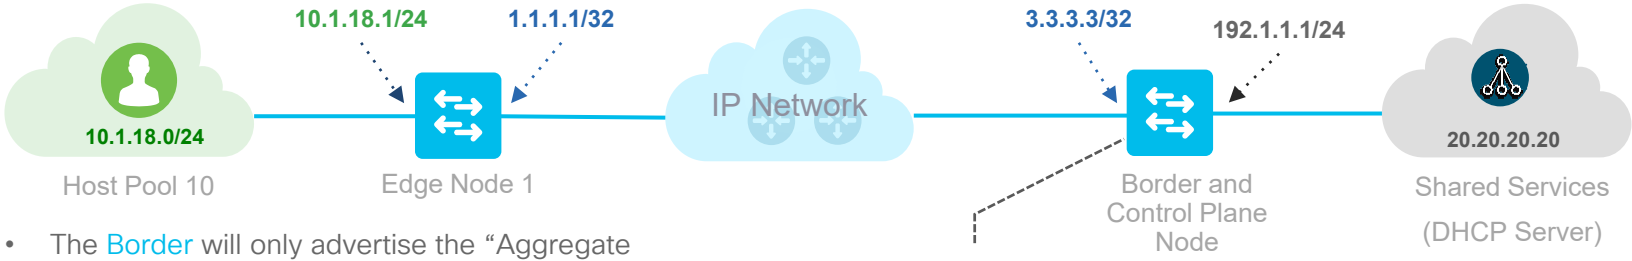

- The Border will only advertise the "Aggregate route(10.1.18.0/24) when there is a more specific prefix(host route)for that subnet in RIB"
- For the DHCP server to send the DHCP reply back we need a route to 10.1.18.0/24 at the DCHP server side.
- This will not happen until we have a HOST in the subnet 10.1.18.0/24 registered with LISP HTDB. This needs DHCP to happen first.

```
router bgp 65002
 bgp log-neighbor-changes
!
address-family ipv4 vrf USER
  aggregate-address 10.1.18.0 255.255.255.0 summary-only
  redistribute lisp metric 10
  neighbor x.x.x.x remote-as xxxxx
  exit-address-family

interface Loopback1022
vrf forwarding User
ip address 10.1.18.1 255.255.255.255
```


© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – Include RLOC in DHCP request

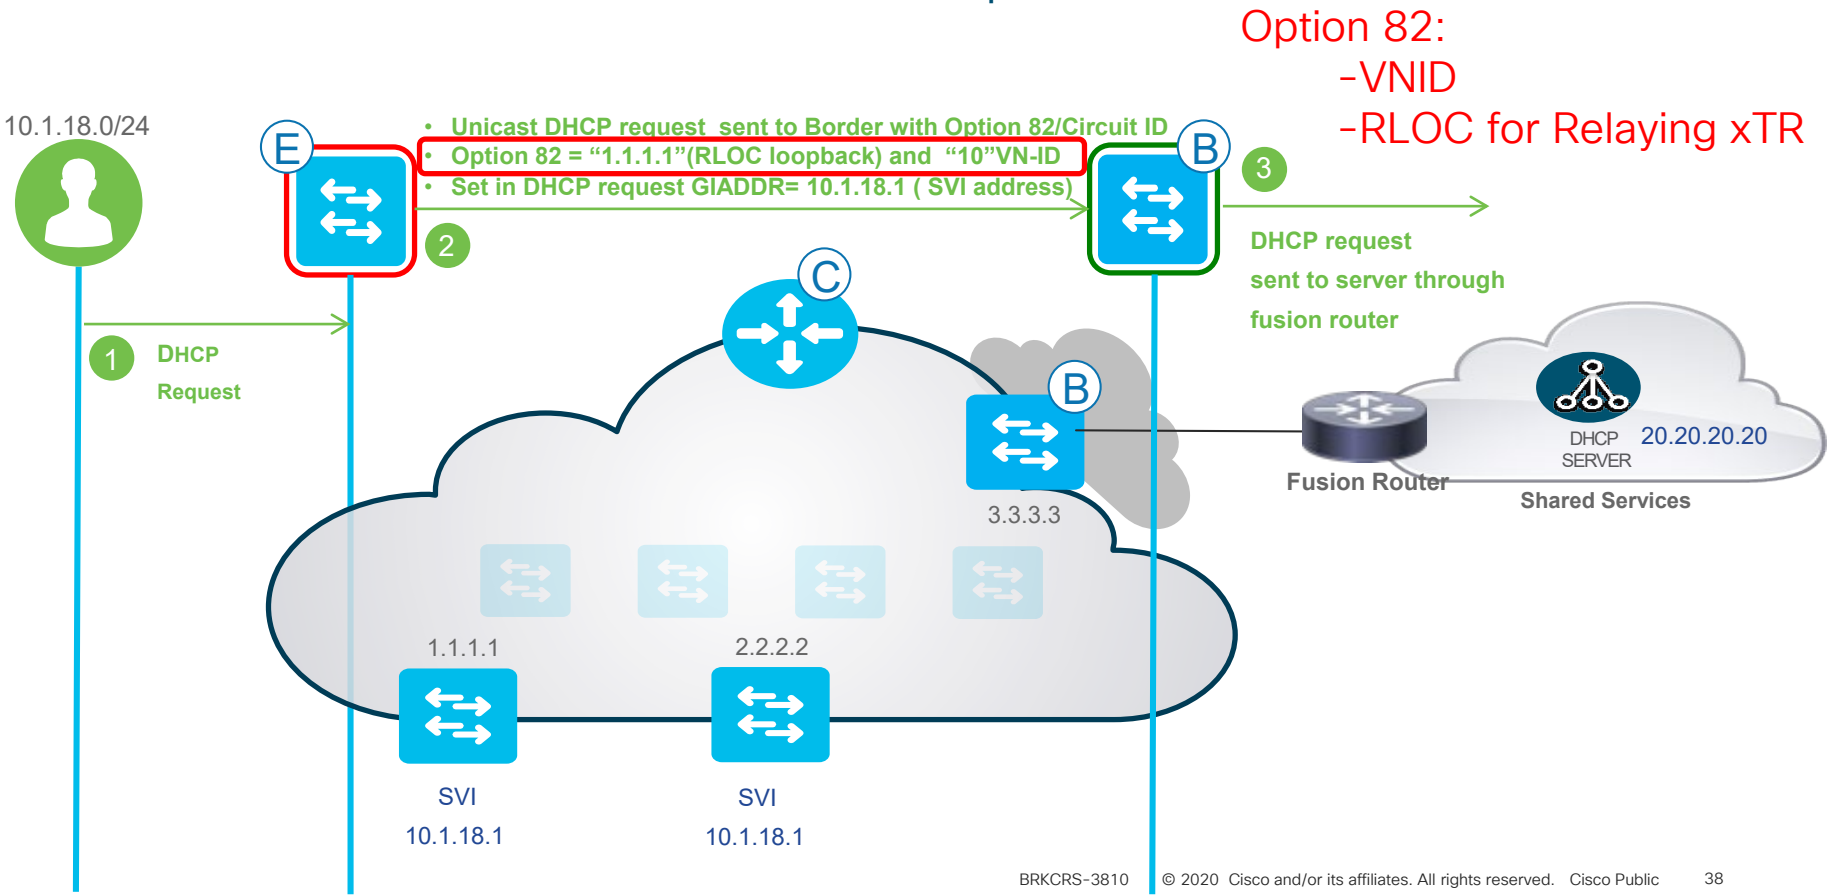

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – DHCP Reply

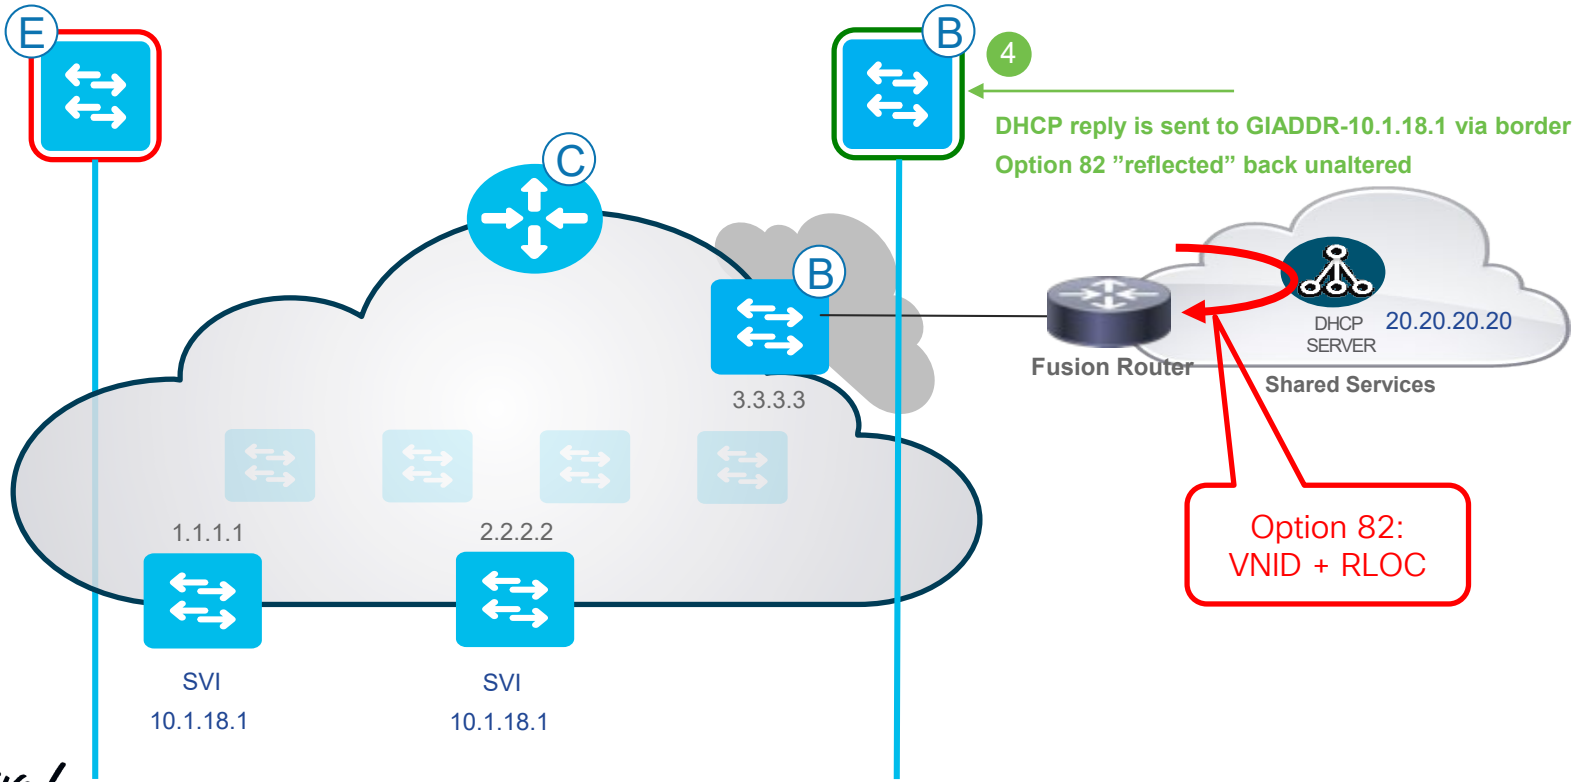

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – Snoop and encapsulate DHCP reply at the Border

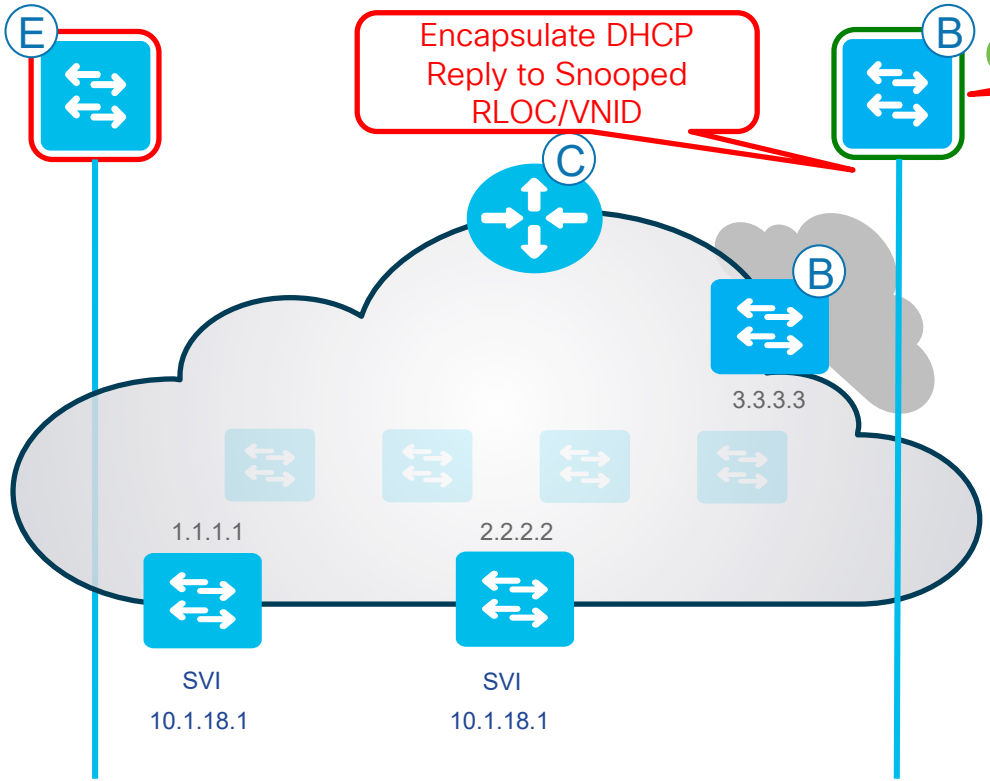

- DHCP reply comes to the Border where it is punted to CPU.
- Relay xTR RLOC + VNI is extracted from Option82 / Circuit ID
- Border encapsulates DHCP reply in VXLAN with destination RLOC = 1.1.1.1 and VNI = 10

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Fabric Architecture

DHCP in Fabric – Complete DHCP reply

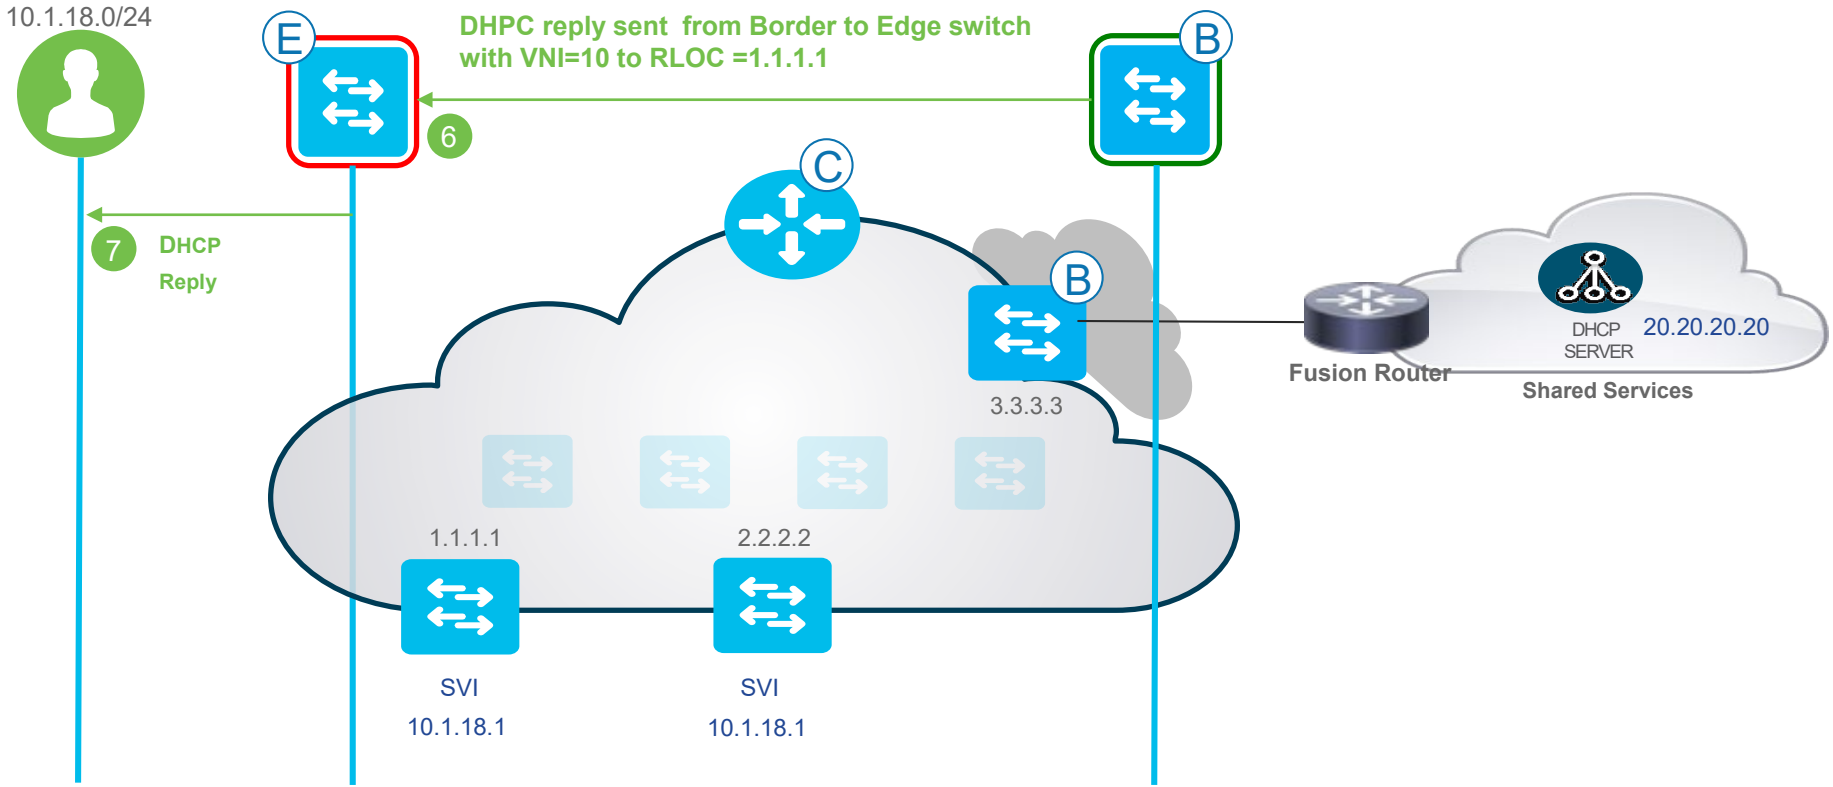

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

1. Control Plane Lookup
2. Access Control Policy

Unicast Packet Forwarding & Access Control

Cisco SD-Access Architecture

Control-Plane Nodes – A Closer Look

RECAP

Control-Plane Node runs a Host Tracking Database to map location information

- A simple Host Database that maps Endpoint IDs to a current Location, along with other attributes
- Host Database supports multiple types of Endpoint ID lookup types (IPv4, IPv6 or MAC)
- Receives Endpoint ID map registrations from Edge and/or Border Nodes for “known” IP prefixes
- Resolves lookup requests from Edge and/or Border Nodes, to locate destination Endpoint IDs

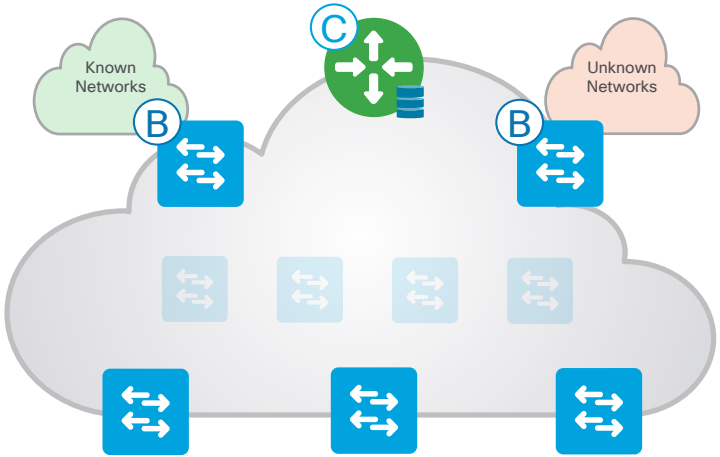

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Endpoint Registration in Fabric

Cisco SD-Access Fabric Architecture

Host Registration

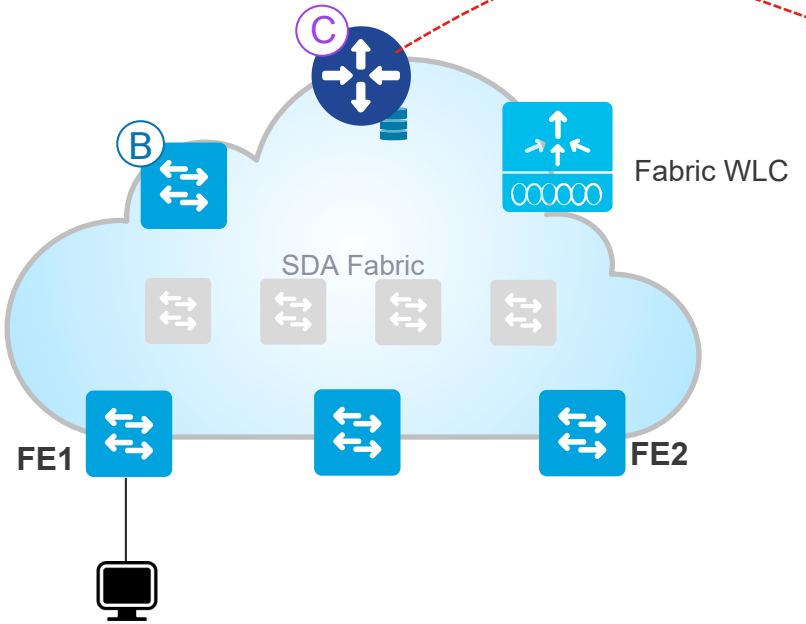

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table

- The Control Plane node has three related tables:
  - **IP to RLOC** — stores IP address of a Host and its corresponding location
  - **MAC to RLOC** — stores MAC address of a Host and its corresponding location
  - **Address Resolution** — Data from the above two tables are collated for IP to MAC bindings (ARP Table)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Host Registration

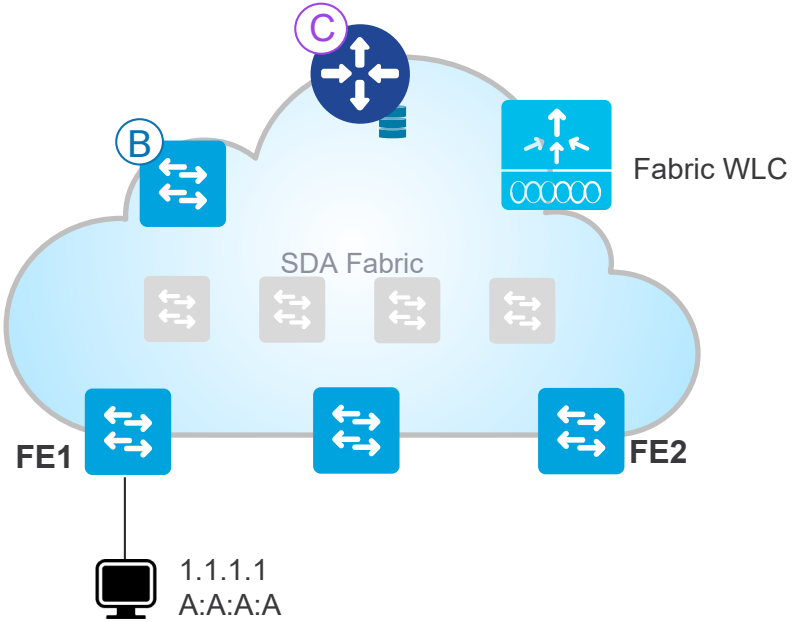

1. Wired host attaches to the fabric network on an edge node

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Host Registration

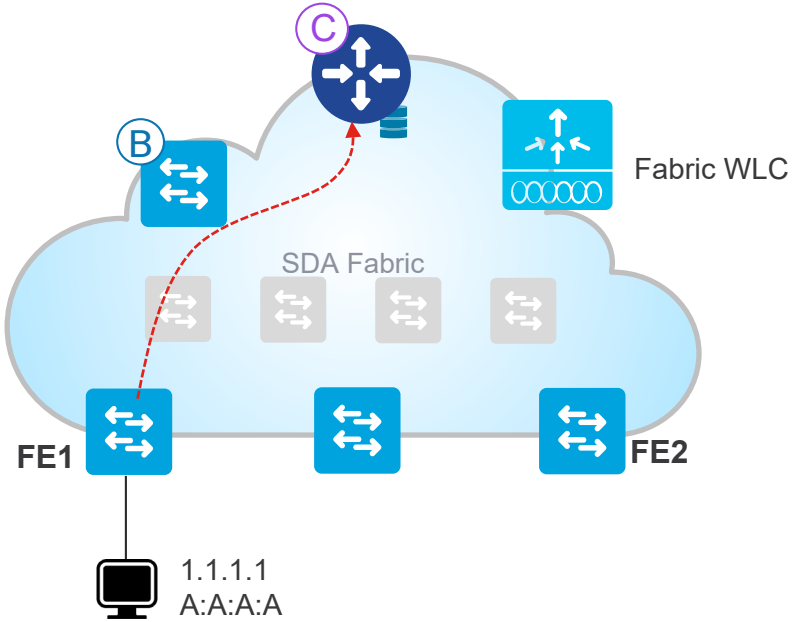

2 After the host gets an IP address, the fabric edge sends a map-register to the control plane node with the IP and mac address of the host

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Host Registration

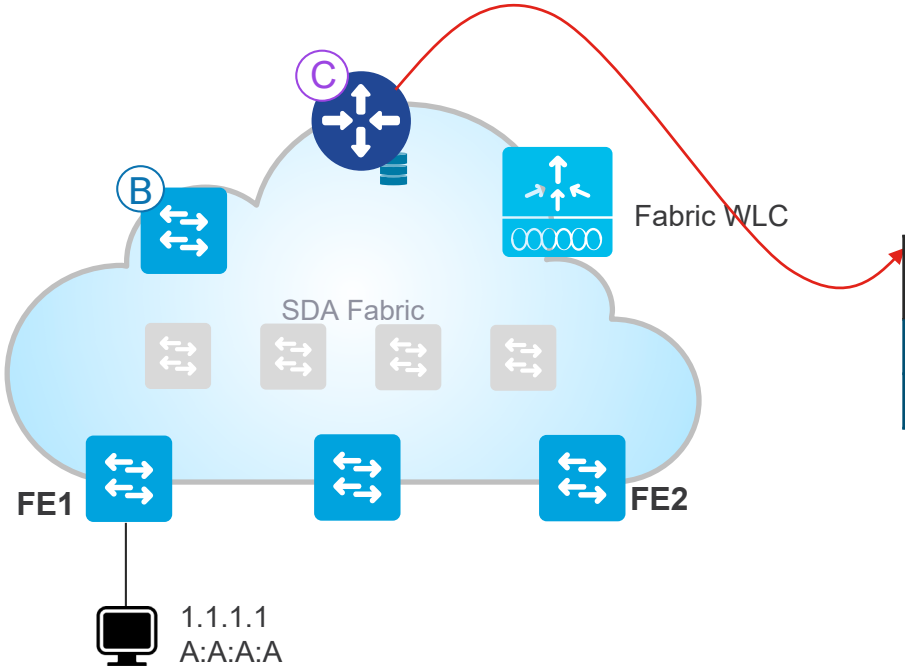

3 The control plane node upon receiving the map-register populates the database tables for the host

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	AA:AA:AA:AA → FE1	

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Host Registration

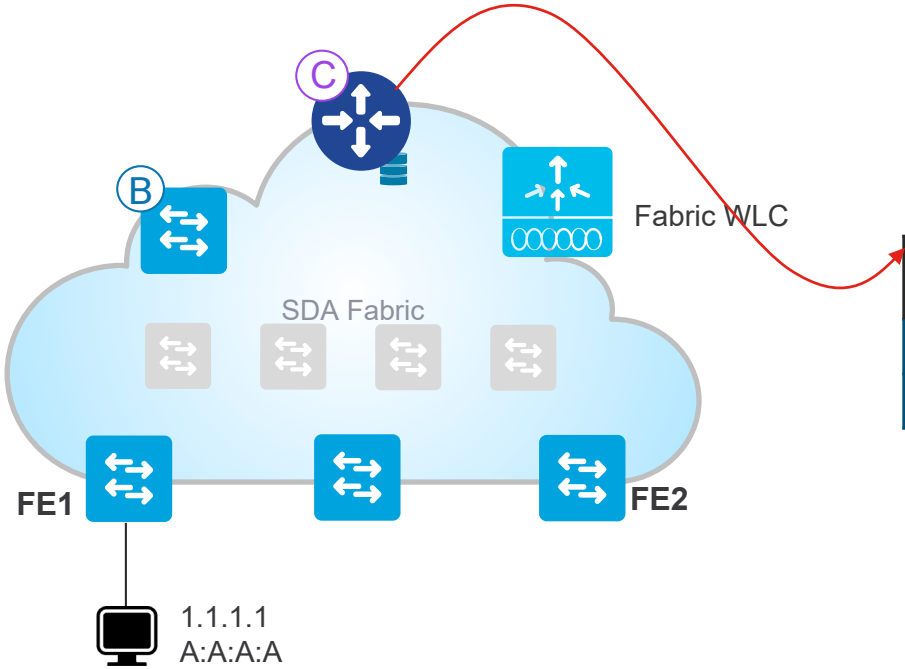

4 The Control plane then takes the information from the IP and MAC table and populates the ARP table

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	AA:AA:AA:AA → FE1	1.1.1.1 = A:A:A:A

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Registering an Access Point in Fabric

SD-Access Wireless Workflow

AP onboarding – AP Discovery and VLAN assignment

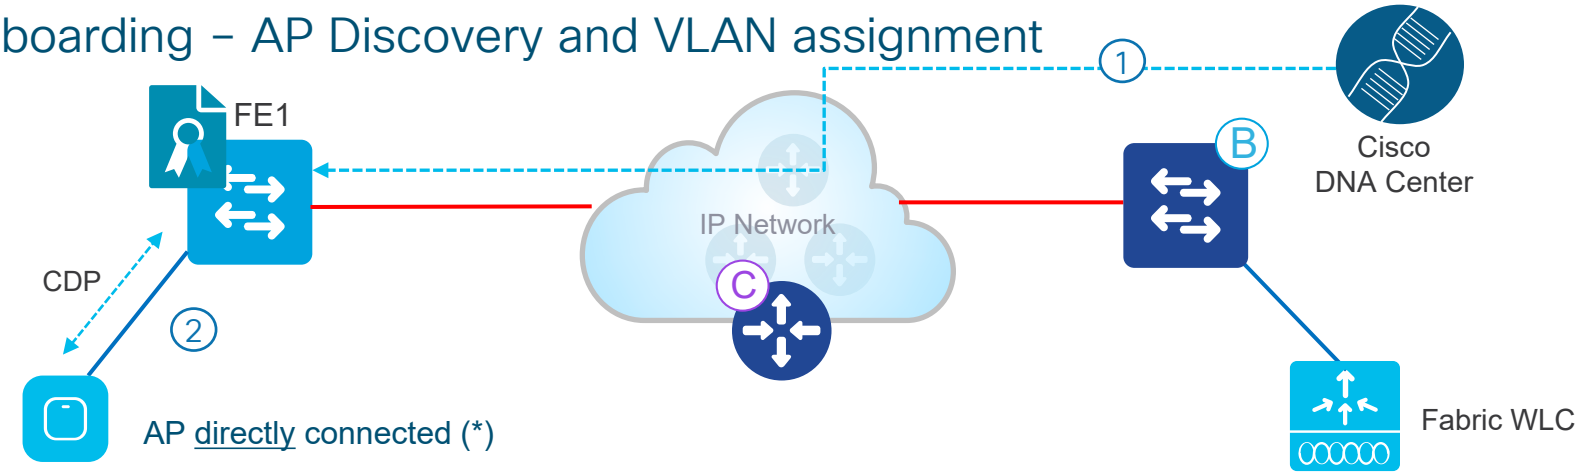

1. Admin configures AP pool in Cisco DNA Center in INFRA_VN. Cisco DNA Center pre-provision a configuration macro on all the FEs
2. AP is plugged in and powers up. FE discovers it's an AP via CDP and applies the macro to assign the switch-port the right VLAN

(*) AP can be connected also through an "Extended node" switch

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

AP onboarding – AP assigned IP address

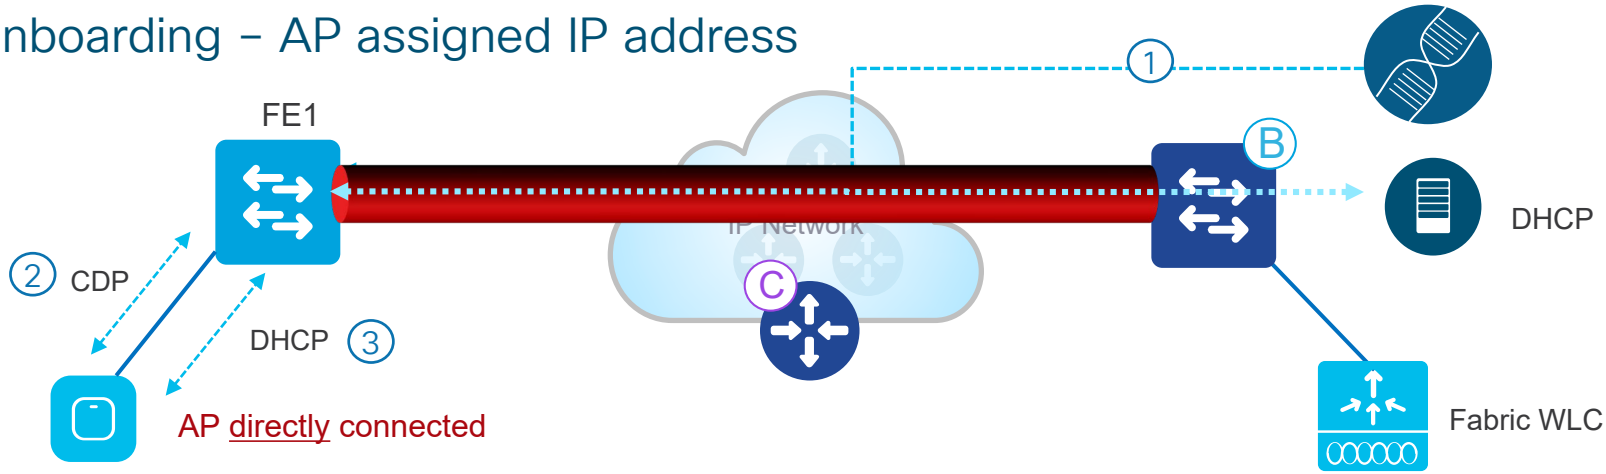

1. Admin configures AP pool in Cisco DNA Center in INFRA_VN. Cisco DNA Center pre-provision a configuration macro on all the FEs
2. AP is plugged in and powers up. FE discovers it's an AP via CDP and applies the macro to assign the switch port the the right VLAN
3. AP gets an IP address via DHCP in the overlay

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

AP onboarding – Register AP with Fabric and Connect to WLC

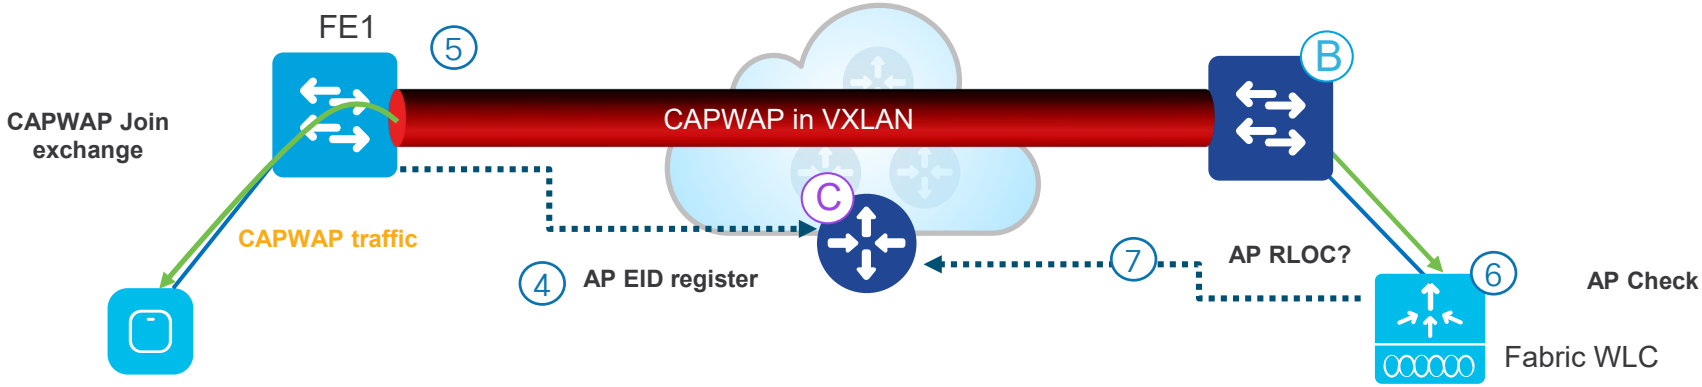

4 Fabric Edge registers AP's IP address and MAC (EID) and updates the Control Plane (CP)

5 AP learns WLC's IP and joins using traditional methods. Fabric AP joins in Local mode

6 WLC checks if AP is fabric-capable (11ax, Wave 2, Wave 1 APs)

7 If AP is supported, WLC queries the CP to know if AP is connected to Fabric

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

AP onboarding – Switch AP to Fabric mode

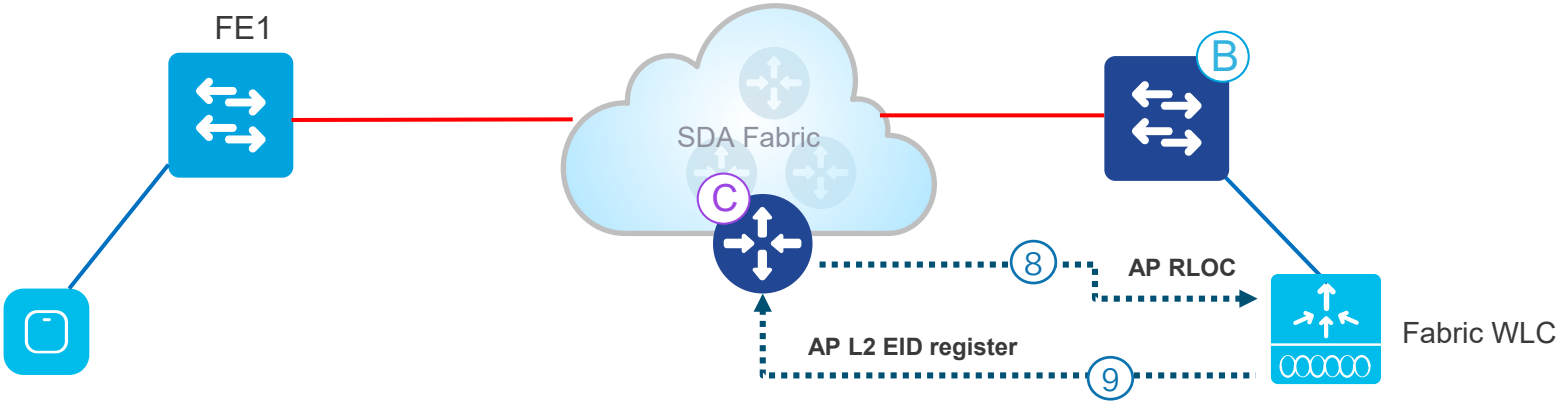

8 Control Plane (CP) replies to WLC with RLOC. This means AP is attached to Fabric and will be shown as “Fabric enabled”

9 WLC does a L2 LISP registration for the AP in CP (a.k.a. AP “special” secure client registration). This is used to pass important metadata information from WLC to the FE

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

AP onboarding – Program access ports for Fabric AP

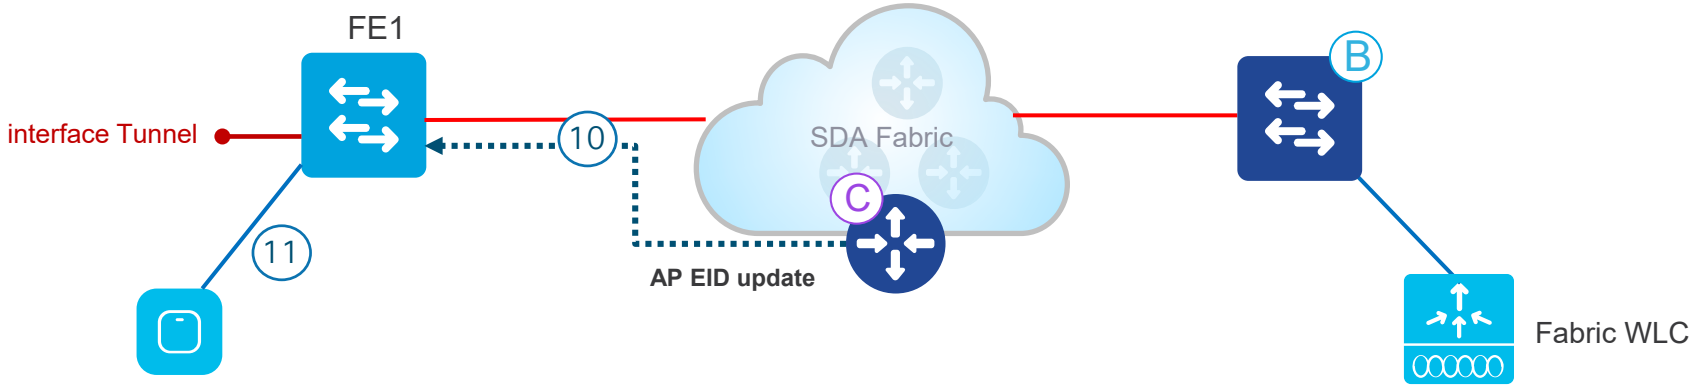

10 In response to this proxy registration, Control Plane (CP) notifies Fabric Edge and pass the metadata received from WLC (flag that says it's an AP and the AP IP address)

11 Fabric Edge processes the information, it learns it's an AP and creates a VXLAN tunnel interface to the specified IP (optimization: switch side is ready for clients to join)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

AP onboarding – Program access ports for Fabric AP

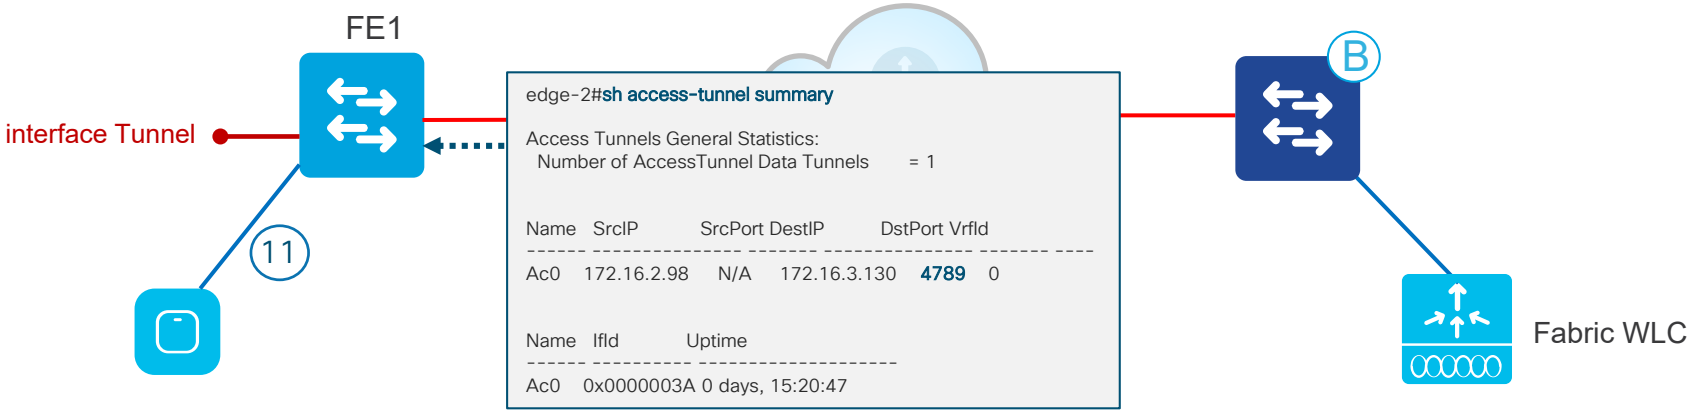

10 In response to this proxy registration, Control Plane (CP) notifies Fabric Edge and pass the metadata received from WLC (flag that says it's an AP and the AP IP address)

11 Fabric Edge processes the information, it learns it's an AP and creates a VXLAN tunnel interface to the specified IP (optimization: switch side is ready for clients to join)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Registering a Wireless host in Fabric

SD-Access Wireless Workflow

Client Onboarding

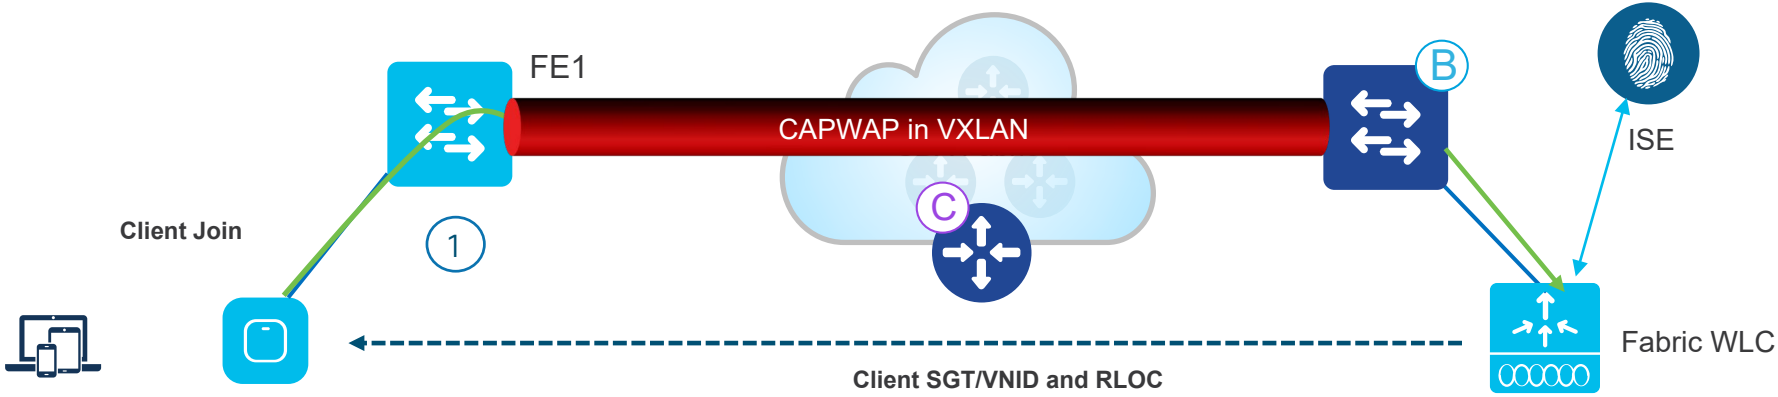

1 Client authenticates to a Fabric enabled WLAN. WLC gets SGT from ISE, updates AP with client L2VNID and SGT

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

Client Onboarding

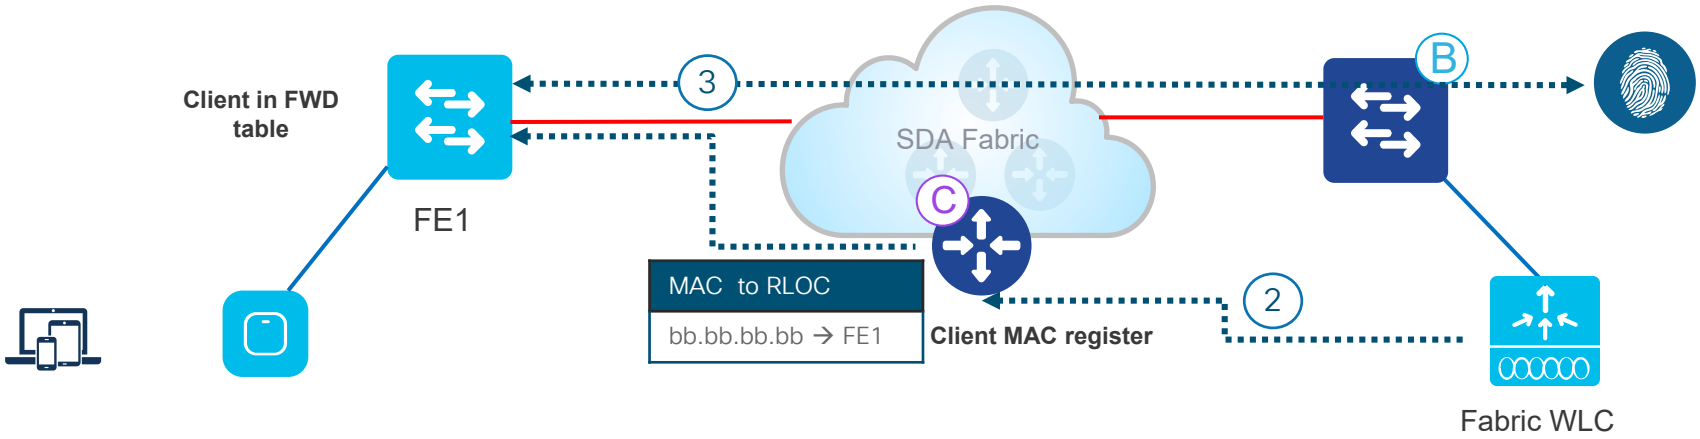

1. Client authenticates to a Fabric enabled WLAN. WLC gets SGT from ISE, updates AP with client L2VNID and SGT
2. WLC knows RLOC of AP from internal DB . WLC proxy registers Client L2 info in CP; this is LISP modified message to pass additional info, like the client SGT
3. FE gets notified by CP and knows it's a client; FE adds client MAC in L2 forwarding table and go and fetch the client policy from ISE based on the client SGT

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

Client Onboarding

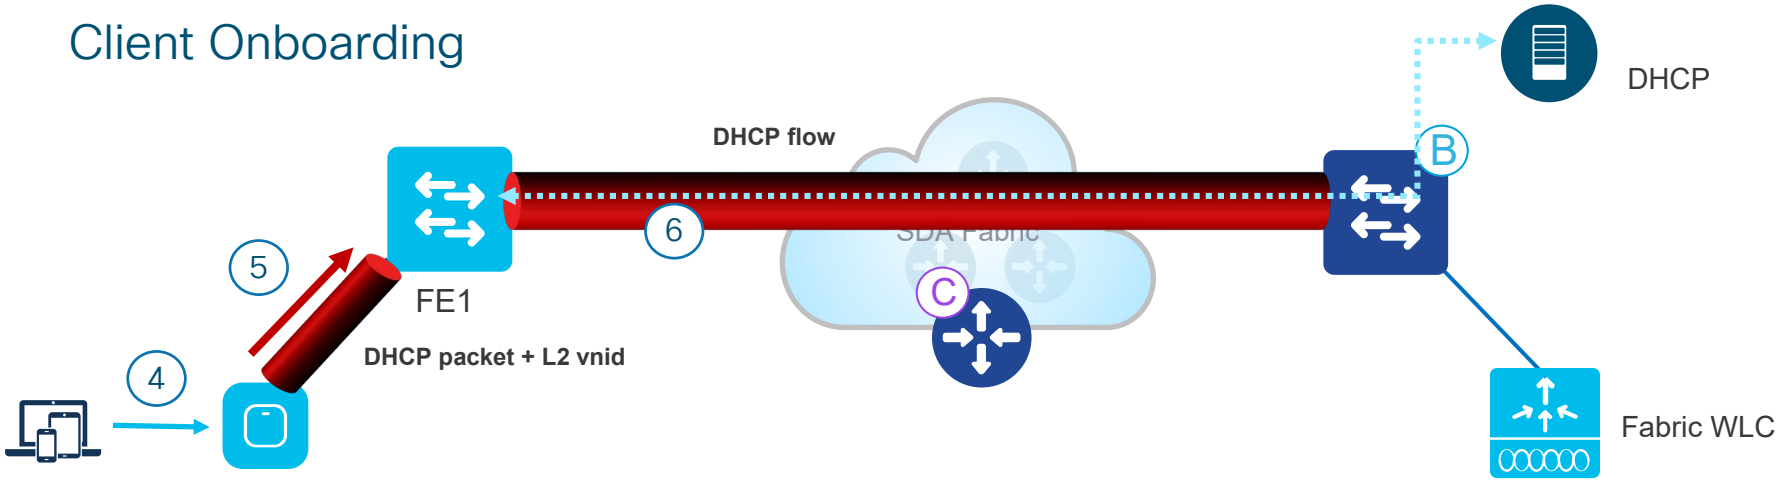

4 Client initiates DHCP Request

5 AP encapsulates it in VXLAN with L2 VNI info (and SGT)

6 Fabric Edge maps L2 VNID to the VLAN interface and forwards the DHCP packet in the overlay (same as for a wired Fabric client)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Wireless Workflow

Client Onboarding

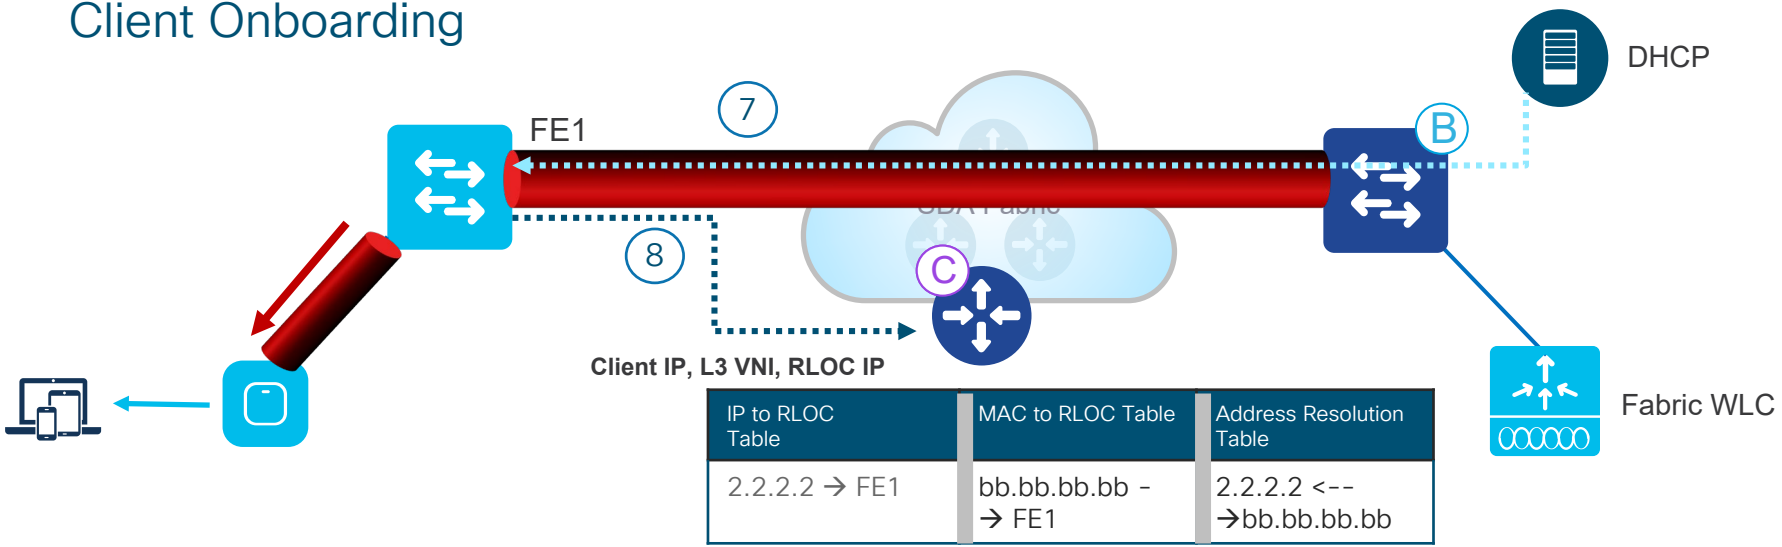

7 Client receives an IP address from DHCP

8 DHCP snooping triggers the client EID registration (MAC, IP address) by the Fabric Edge to the CP. *(If client has a static IP, then ARP or any other IP packet will trigger the registration)*

This completes client onboarding process

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Mobility Events in Fabric

Cisco SD-Access Fabric Architecture

Client Roam

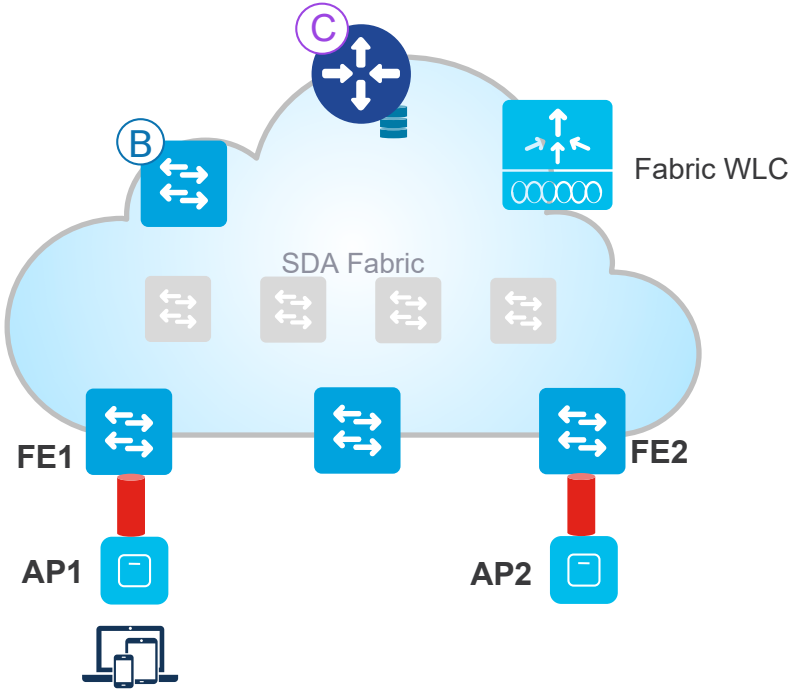

1. Client roams from AP1 to AP2 (inter-switch roaming)

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

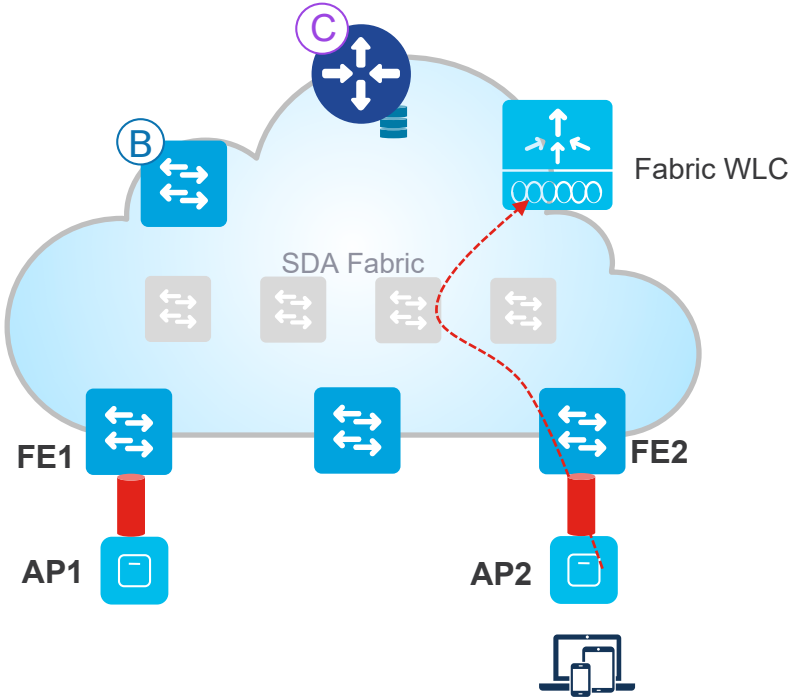

2 AP2 detects the client move and registers its MAC address to the WLC as a mobility event

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

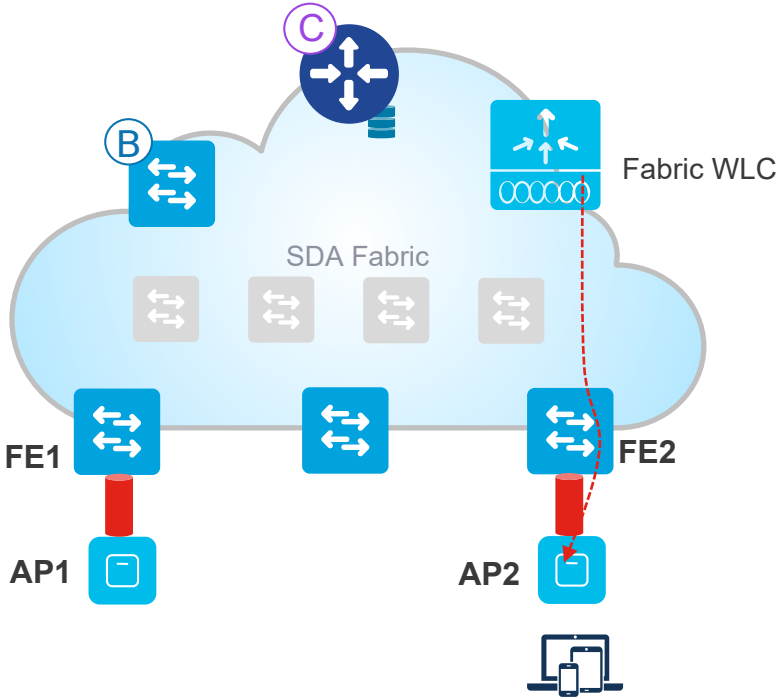

3 WLC updates forwarding table on AP2 with client info (SGT, L2VNID, RLOC)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

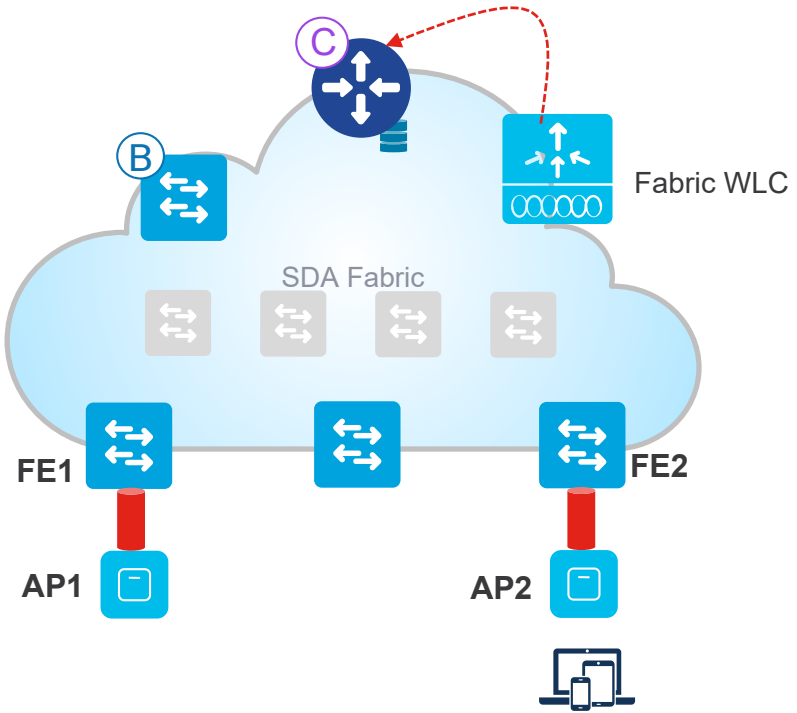

4 WLC updates the L2 MAC entry in CP with new RLOC FE2.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

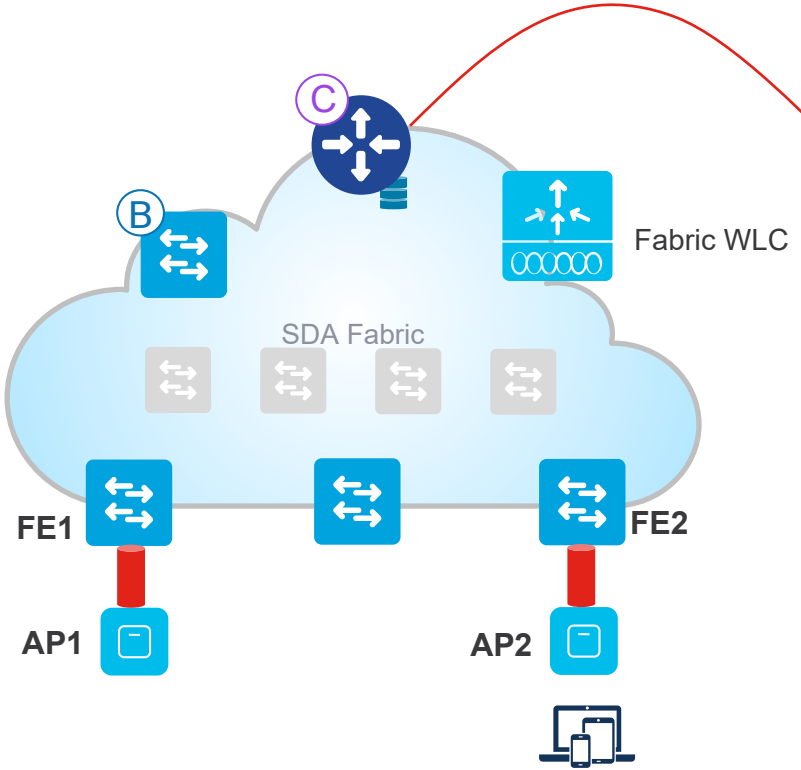

5 The control plane state will be mismatched as the IP address of host says it is located behind FE1 but MAC address says its behind FE2

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
2.2.2.2 → FE1	bb.bb.bb.bb -→ FE2	2.2.2.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

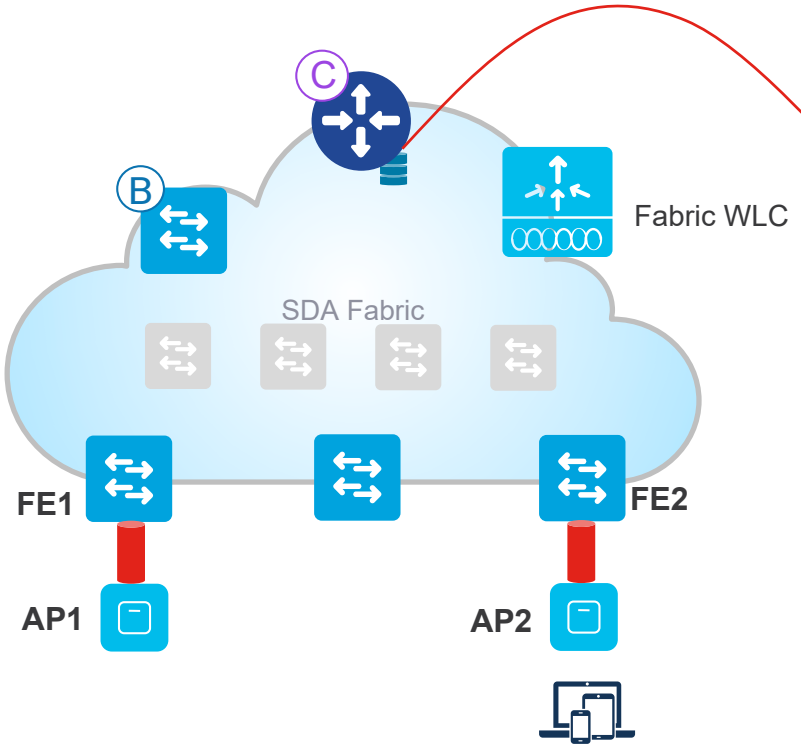

6 In this state the control plane node will trust the notification from the WLC and will correct its IP to RLOC table.

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
2.2.2.2 → FE2	bb.bb.bb.bb - → FE2	2.2.2.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Client Roam

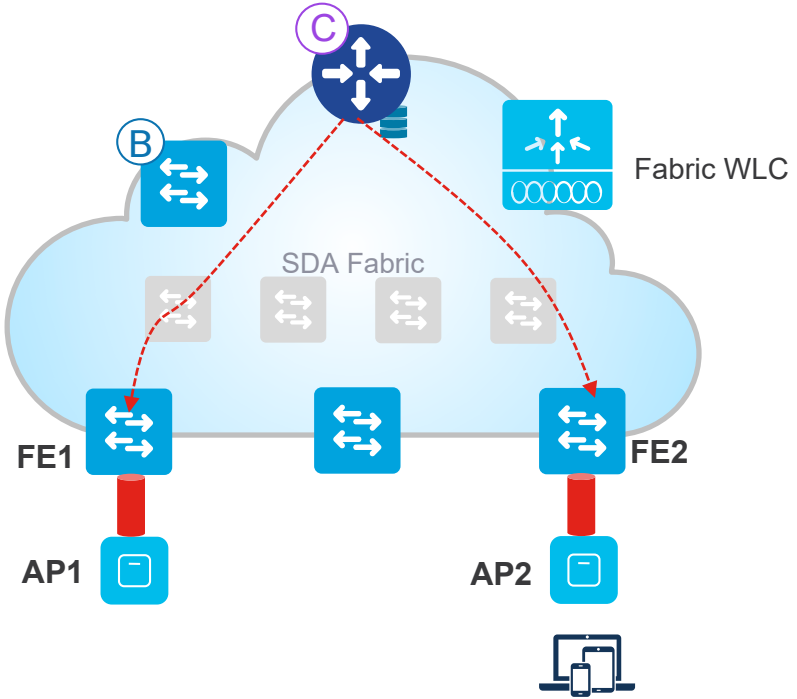

7 The control plane node after correcting its IP to RLOC table will send notifications to Edge 1 and Edge 2.
CP then notifies

- Fabric Edge FE2 (”roam-to” switch) to add the client MAC to forwarding table pointing to VXLAN tunnel
- Fabric Edge FE1 (”roam-from” switch) to do clean up for the wireless client

This ensures that when a host moves from one AP to other, the new edge node is waiting for it even before the move fully completes to ensure seamless roaming.

8 Roam is Layer 2 as FE2 has the same VLAN interface as FE1 (Anycast Gateway)

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Unicast Packet Forwarding in Fabric in the same subnet

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

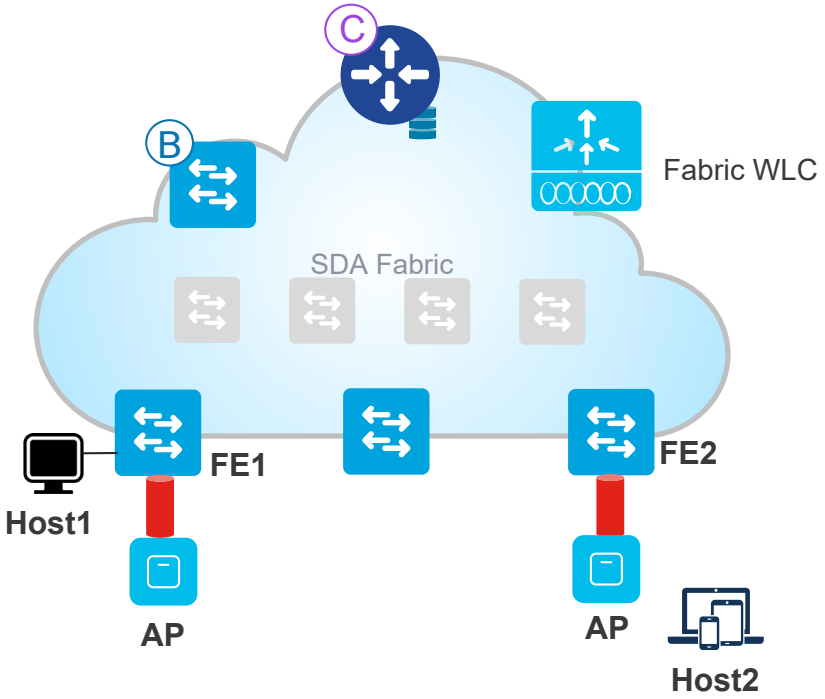

Assumptions:

- Host H1 and Host H2 belong to the same subnet

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa -→ FE1	1.1.1.1 <--→aa.aa.aa.aa
1.1.1.2 → FE2	bb.bb.bb.bb -→ FE2	1.1.1.2 <--→bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

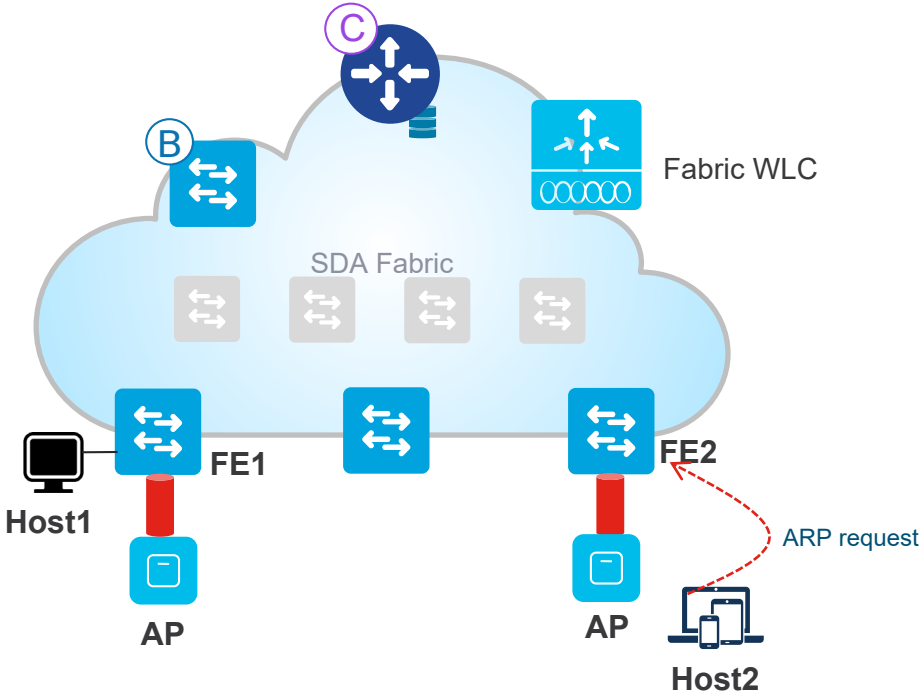

1 Host2 wants to communicate to Host1. Since they are in the same subnet Host2 sends an ARP request for the mac-address of Host1

The packet will contain below:

ARP REQUEST:

SRC IP: H2 IP
DST IP: H1 IP
SRC MAC: H2 MAC
DST MAC: FF:FF:FF:FF

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

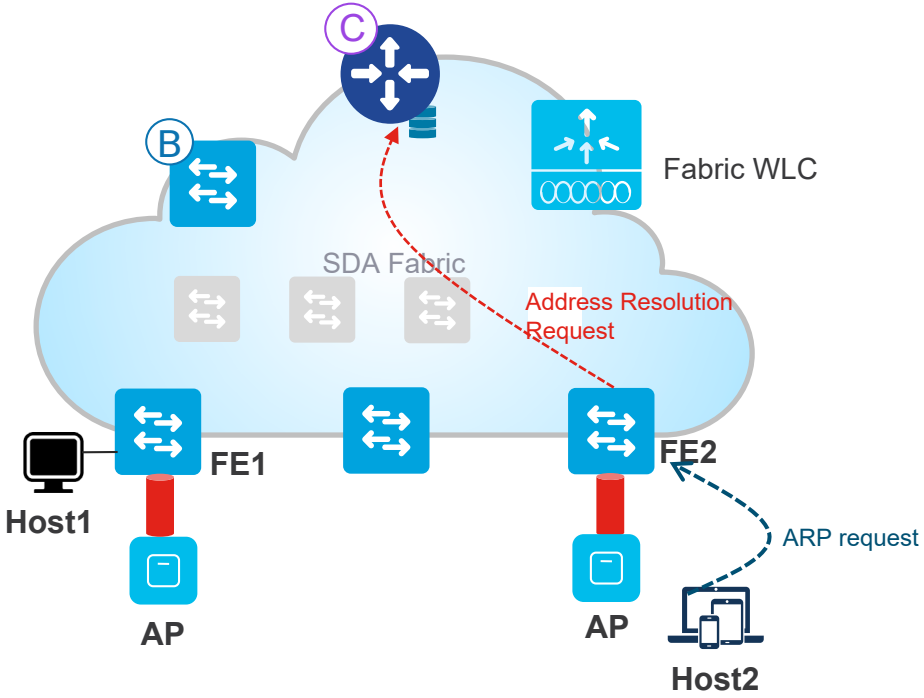

2 Edge node FE2 will intercept the ARP request from Host2 and then will contact the control plane to ask for the MAC address of Host1.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

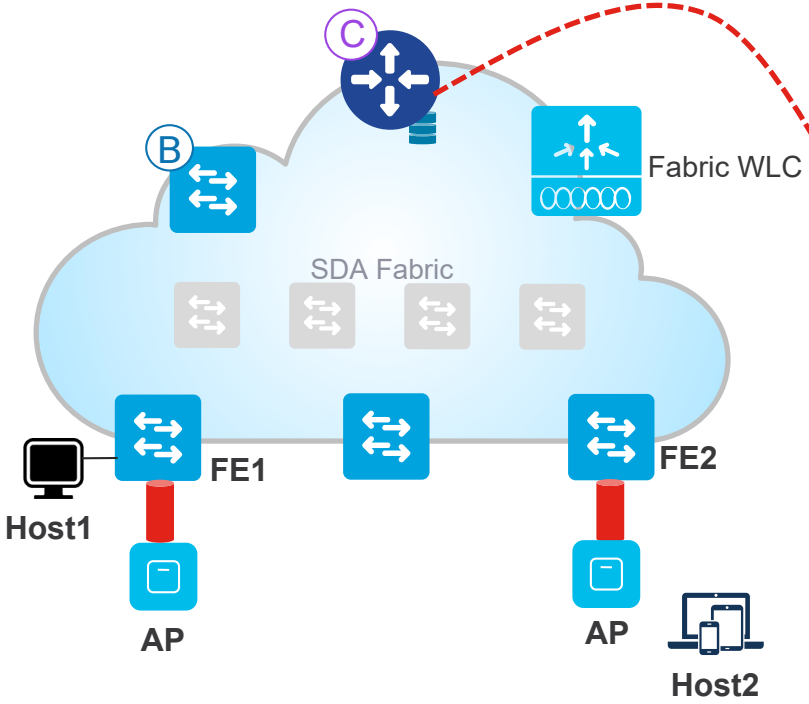

3 The control plane node will consult its Address resolution table to find the IP to MAC binding for Host1

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa - → FE1	1.1.1.1 <-- →aa.aa.aa.aa
1.1.1.2 → FE2	bb.bb.bb.bb - → FE2	1.1.1.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

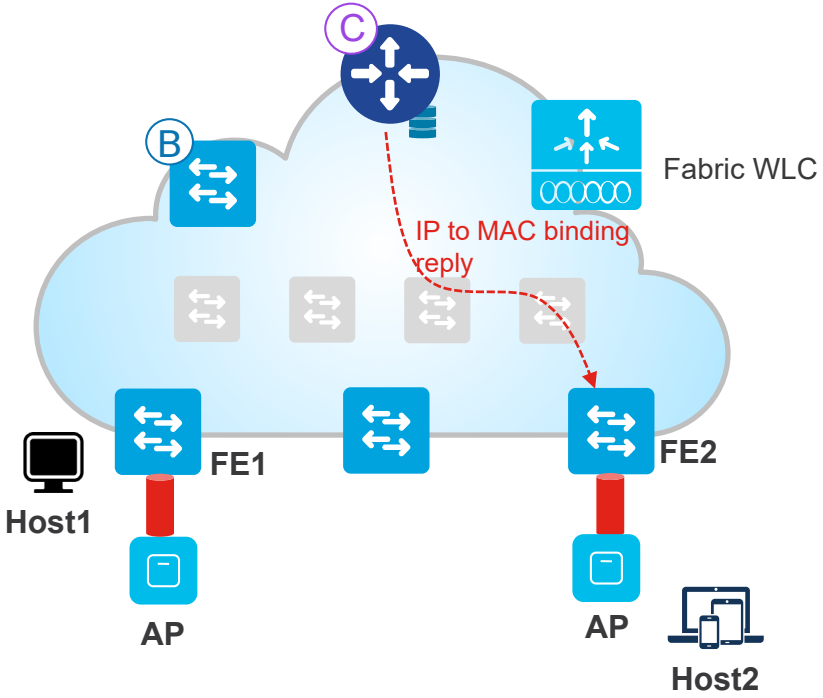

4 If it finds the mapping it will reply back to the edge node with the MAC address of Host1.

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa - → FE1	1.1.1.1 <-- →aa.aa.aa.aa
1.1.1.2 → FE2	bb.bb.bb.bb - → FE2	1.1.1.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

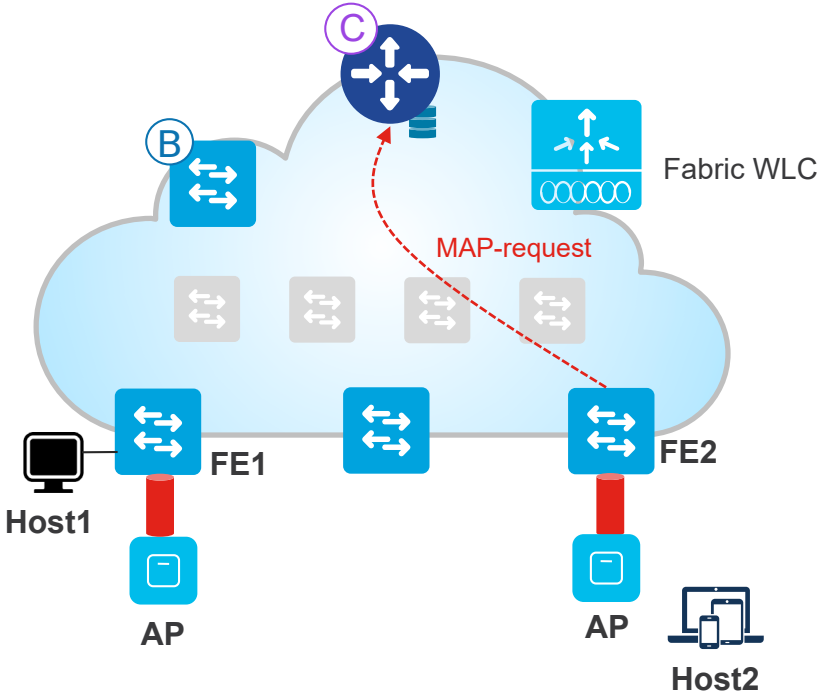

5 Once the edge node FE2 gets the MAC address for Host1, it will again consult the control plane node to find out the location of Host1's MAC address.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

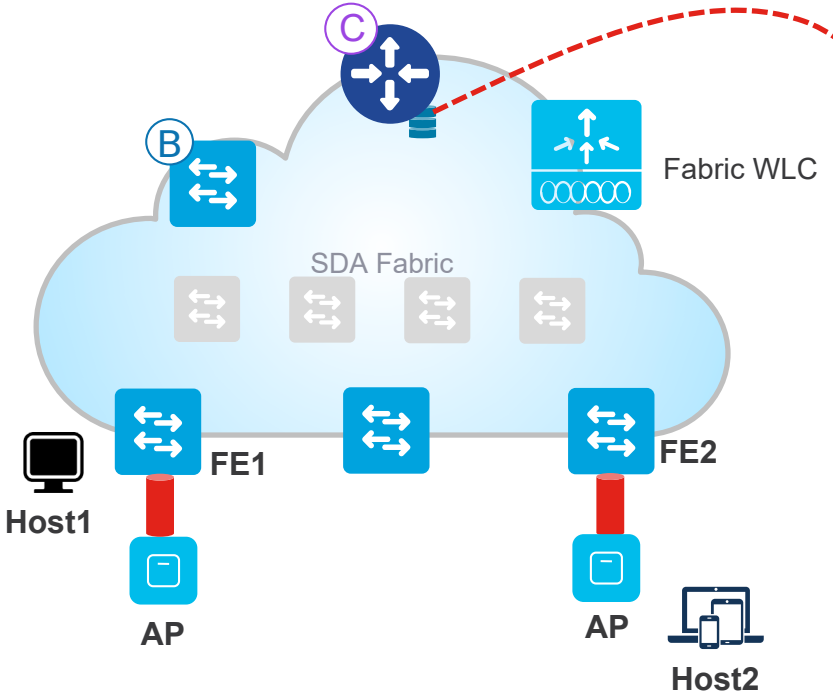

6 The control plane node after getting the map-request will consult its MAC to RLOC table for the location.

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa → FE1	1.1.1.1 <-- →aa.aa.aa.aa
1.1.1.2 → FE2	bb.bb.bb.bb → FE2	1.1.1.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

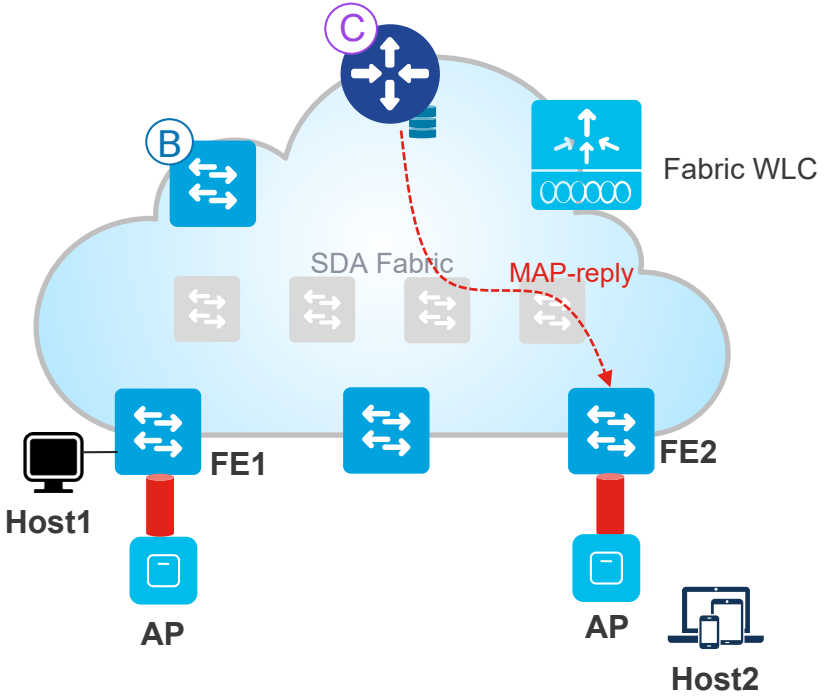

7 If the entry is found the information is sent to the edge node

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa - → FE1	1.1.1.1 <-- →aa.aa.aa.aa
1.1.1.2 → FE2	bb.bb.bb.bb - → FE2	1.1.1.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

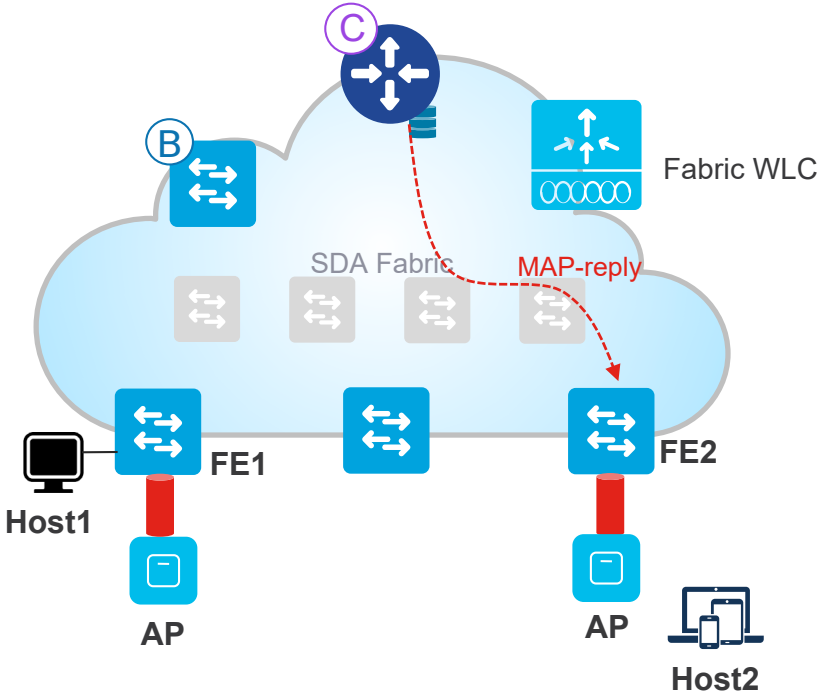

8 The fabric edge node upon receiving the map-reply will install this entry into the Layer 2 forwarding table.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

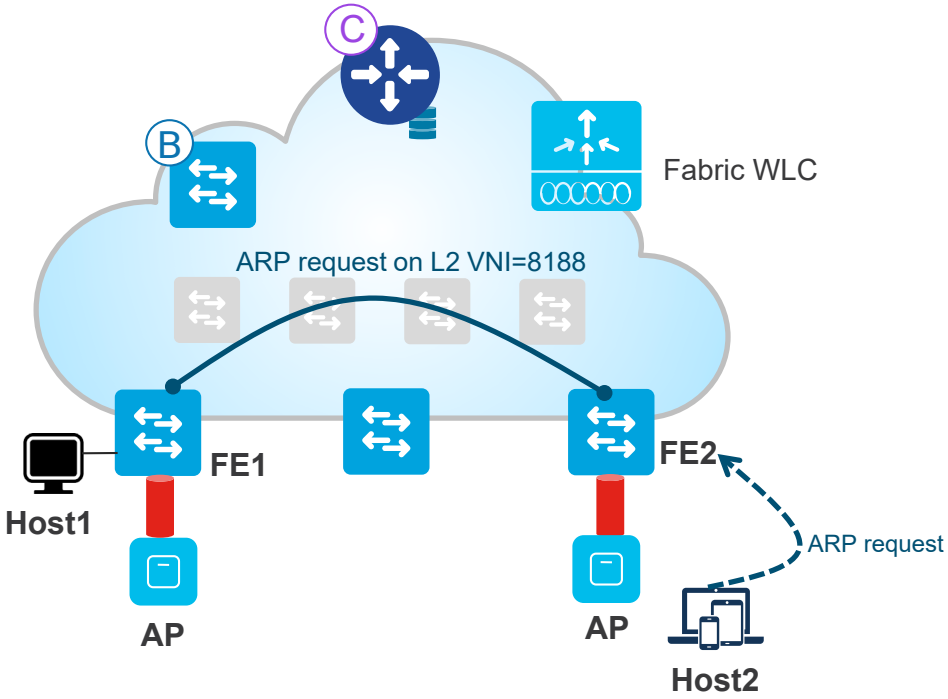

9 Edge node FE2 will now convert the ARP broadcast it received from Host2 to a directed unicast. It will send it to FE1 in the overlay encapsulating it with the L2VNI for that subnet/vlan.

The L2 VNI is derived from the L2 LISP configuration:

```
instance-id 8188
 service ethernet
  eid-table vlan 1024
```

The packet will contain below:

ARP REQUEST:

SRC IP: H2 IP
DST IP: H1 IP
SRC MAC: H1 MAC
DST MAC: H2 MAC << unicast ARP>>>

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

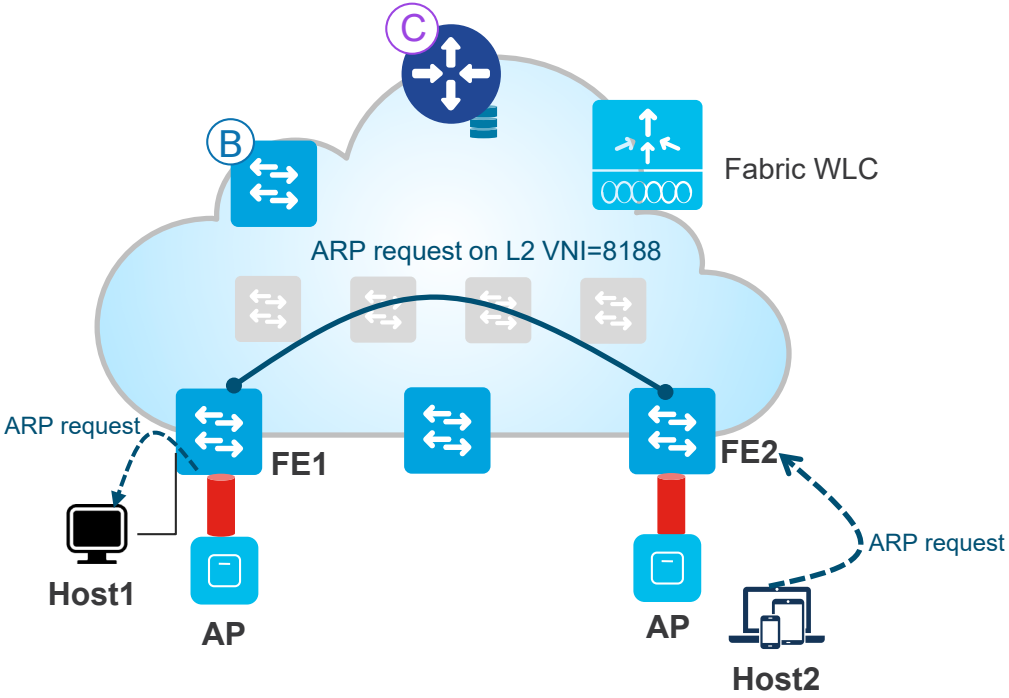

10 The edge node FE1 upon receiving the ARP request will decapsulate the packet. It will forward it to the destination host Host1.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

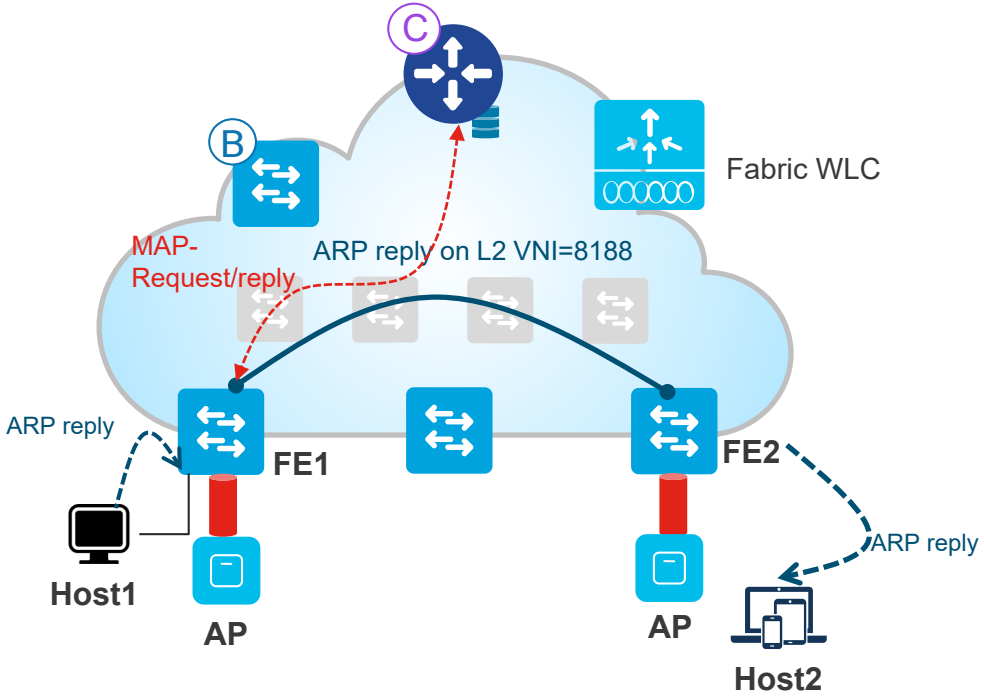

11 The destination Host1 will unicast the ARP reply to Host2.
To do so, FE1 will issue a Map-request for the MAC address of Host 2 and cache the mapping received in the map-reply.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

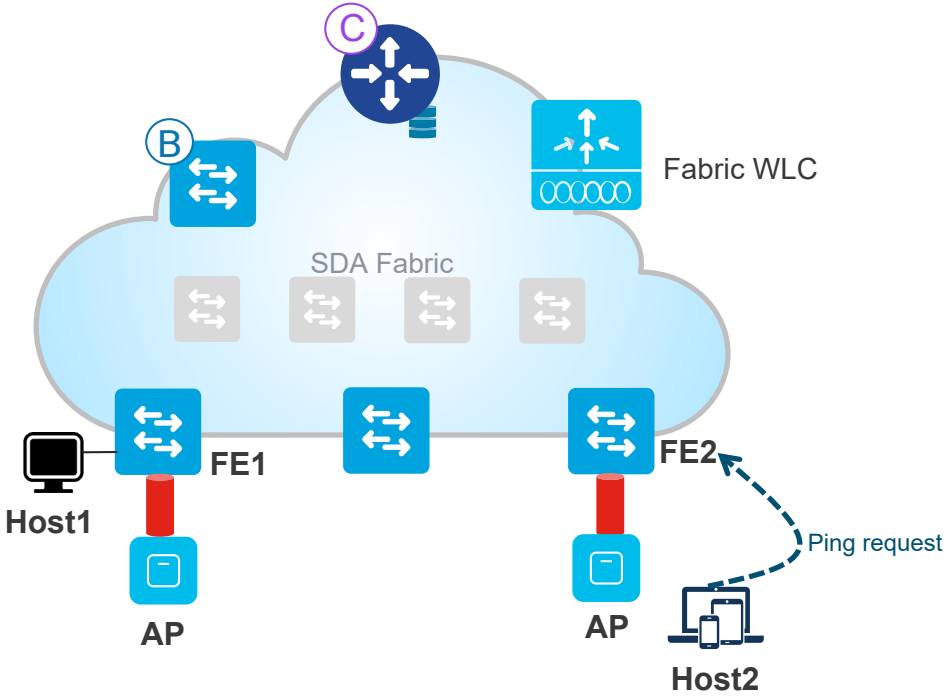

12 Now that ARP is resolved, we can test data connection (ping request) .

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

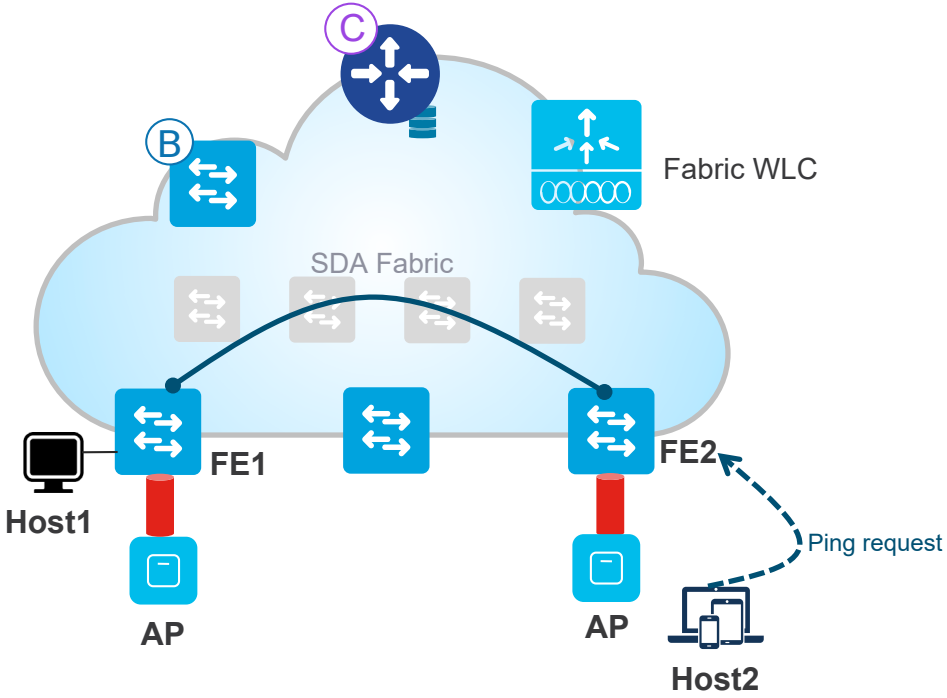

13 The fabric edge node will use the Layer 2 forwarding tables that have been already populated and will use the L2VNI to forward to the destination node.

INNER PACKET:

SRC IP: H2 IP
DST IP: H1 IP
SRC MAC: H2 MAC
DST MAC: H1 MAC

OUTER PACKET:

SRC IP: Edge 2 IP
SRC MAC: Edge 2 MAC
DST IP: Edge 1 IP
DST MAC: Next-Hop MAC (Intermediate node)

.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding in the same subnet

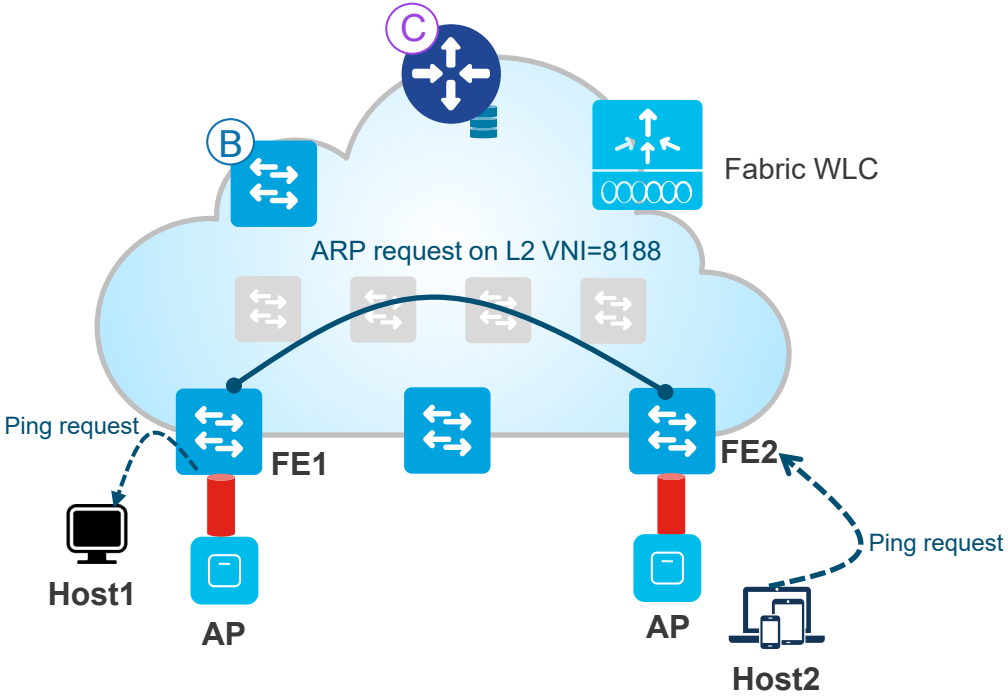

14 Once traffic reaches edge 1 it will decapsulate the VXLAN headers and based on the Inner packet details will do a switching lookup to send it to Host1.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Forwarding across different subnets

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

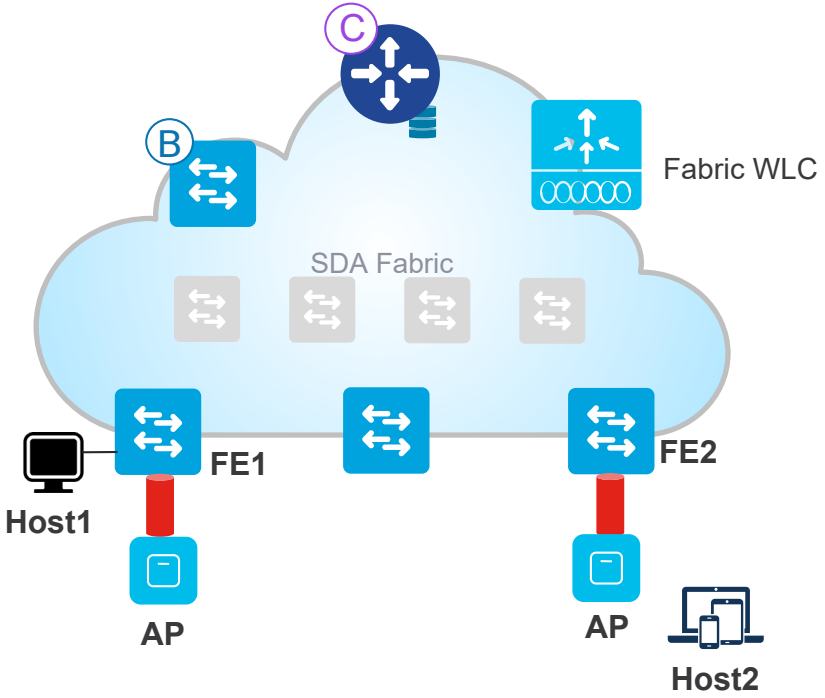

Assumptions:

- Host H1 and Host H2 belong to different subnets

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa - → FE1	1.1.1.1 <-- →aa.aa.aa.aa
2.2.2.2 → FE2	bb.bb.bb.bb - → FE2	2.2.2.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

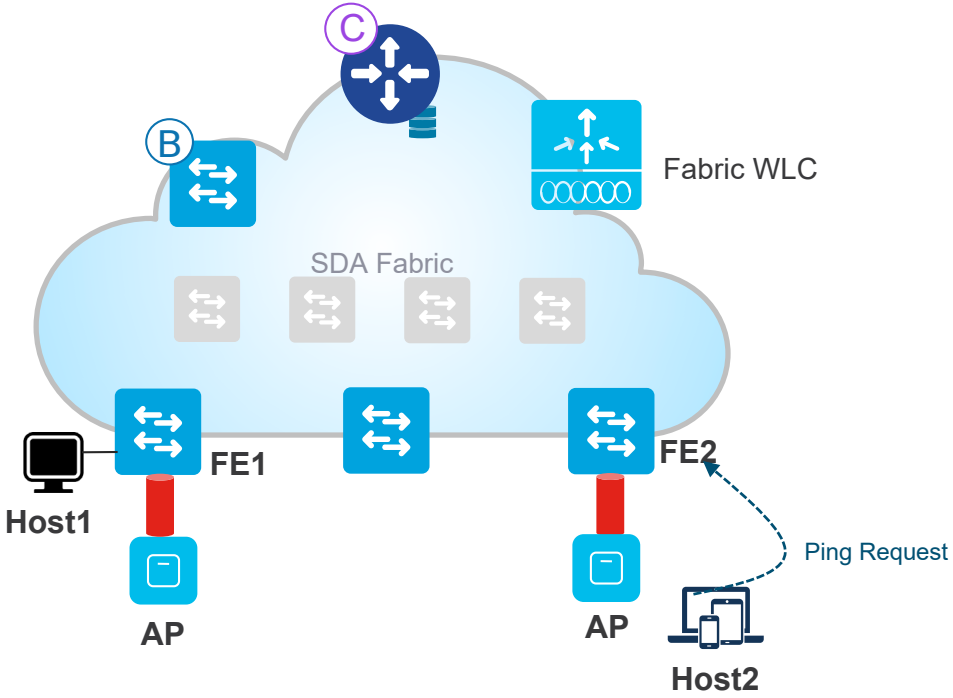

1 Host2 sends a ping packet to Host1.

The packet will contain below:

PING REQUEST:

SRC IP: H2 IP
DST IP: H1 IP
SRC MAC: H2 MAC
DST MAC: DEFAULT GATEWAY MAC

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

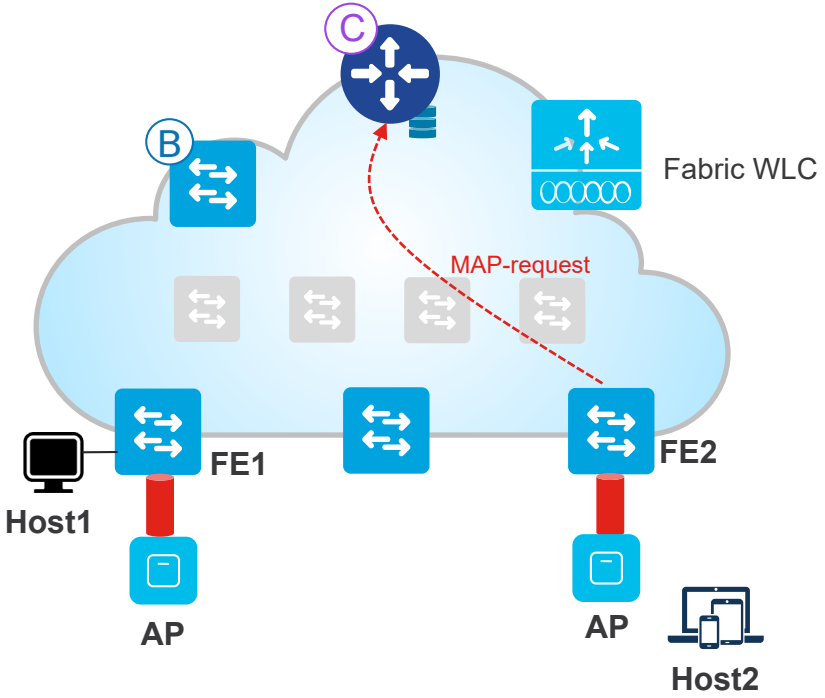

2 Edge node FE2 will intercept the ping request from Host2 destined to Host1 and then will contact the control plane to ask for the location of Host1.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

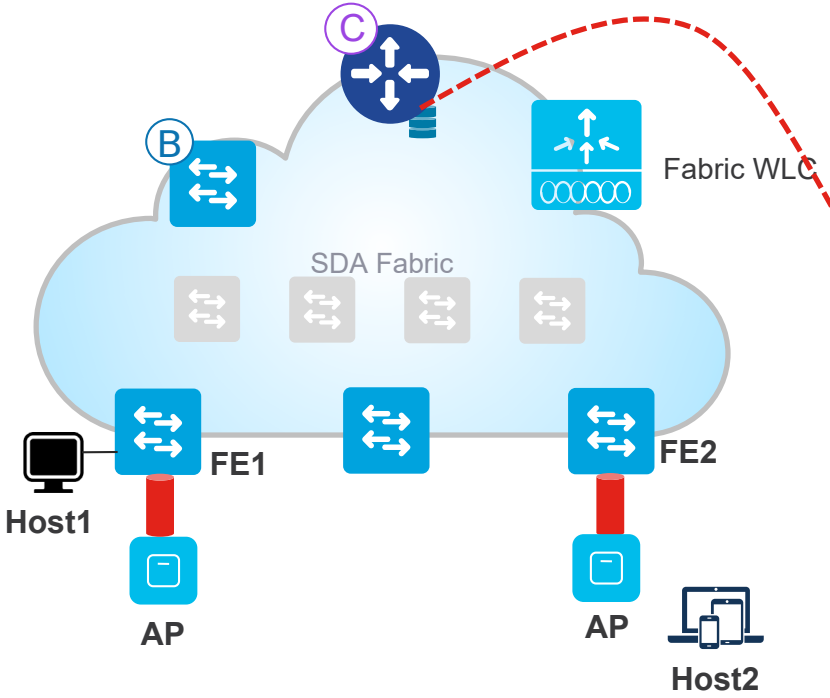

3 The control plane node will consult its IP to RLOC binding for location of Host1.

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa - → FE1	1.1.1.1 <-- →aa.aa.aa.aa
2.2.2.2 → FE2	bb.bb.bb.bb - → FE2	2.2.2.2 <-- →bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

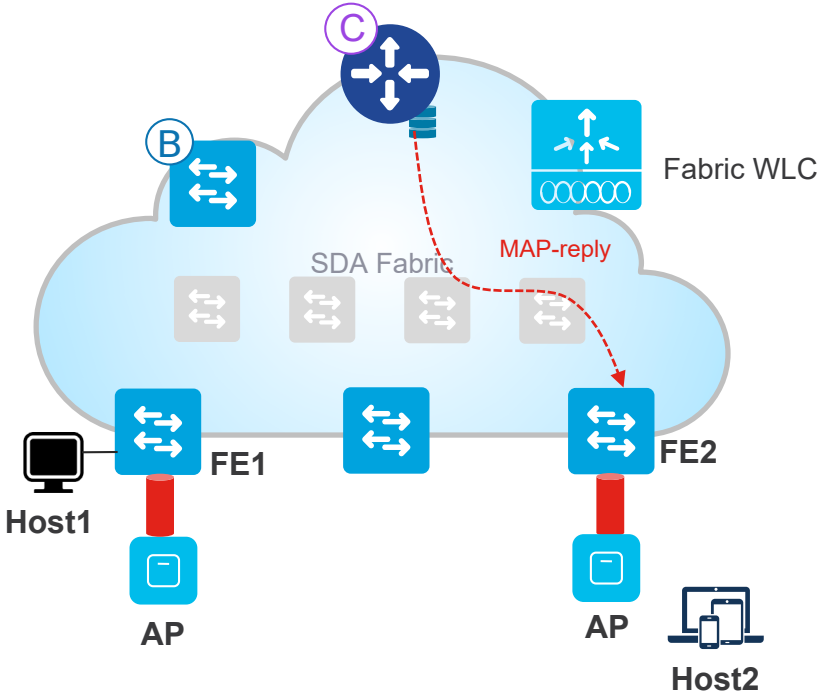

4 If it finds the mapping it will reply back to the edge node with the location of Host1

Control Plane state:

IP to RLOC Table	MAC to RLOC Table	Address Resolution Table
1.1.1.1 → FE1	aa.aa.aa.aa -→ FE1	1.1.1.1 <--→aa.aa.aa.aa
2.2.2.2 → FE2	bb.bb.bb.bb -→ FE2	2.2.2.2 <--→bb.bb.bb.bb

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

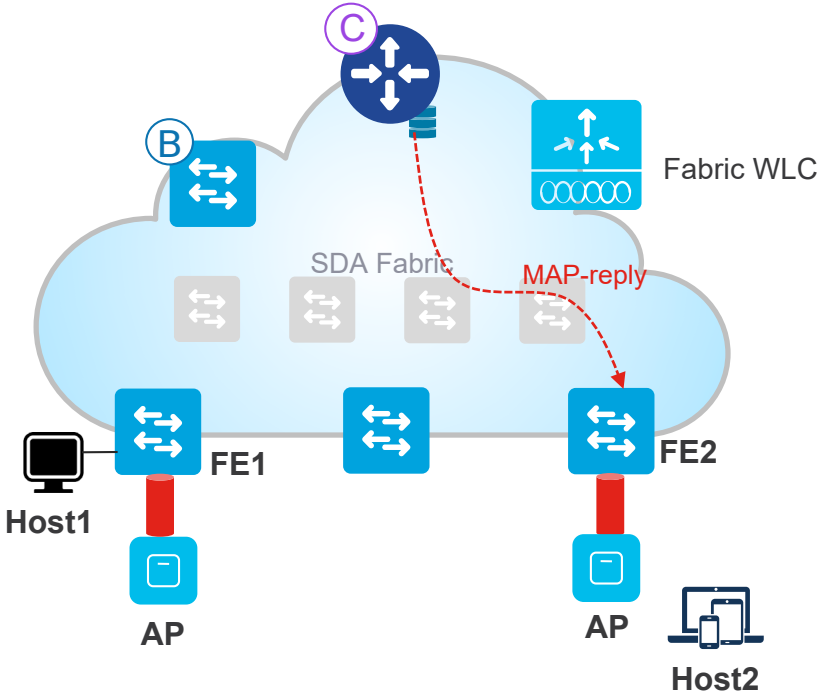

5 The fabric edge node upon receiving the map-reply will install this entry into the Layer 3 forwarding tables .

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

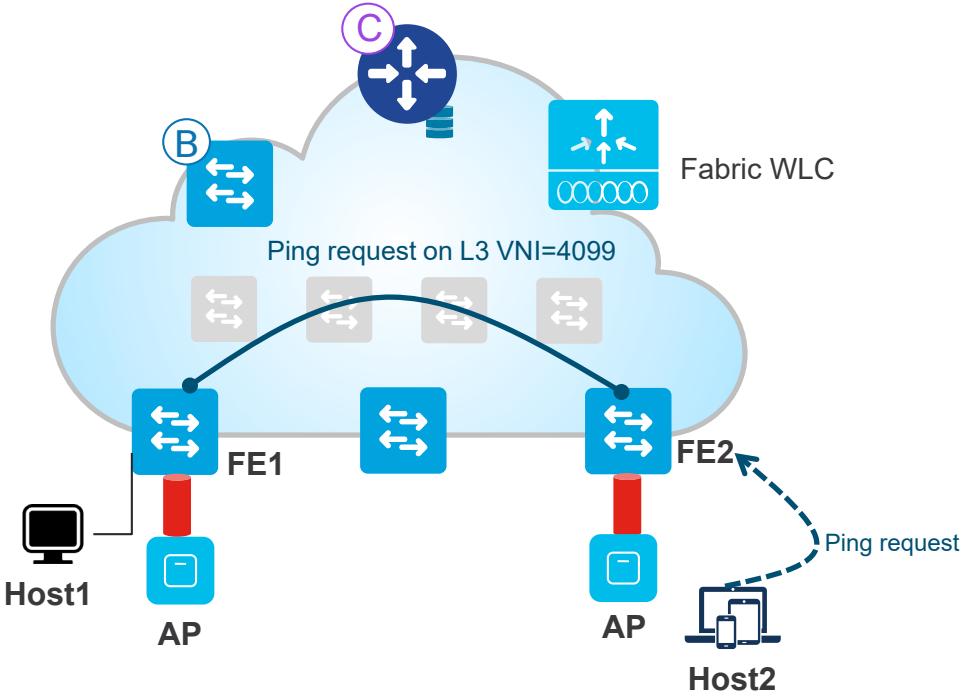

6 Edge node FE2 will now send the ping request received from Host2 to edge node FE1 in the overlay encapsulating it with the L3VNI for that VRF.

The L3 VNI is derived from the LISP configuration:

```
instance-id 4099
 remote-rloc-probe on-route-change
 service ipv4
  eid-table vrf Campus
  map-cache 0.0.0.0/0 map-request
  exit-service-ipv4
```


© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Unicast Forwarding across different subnets

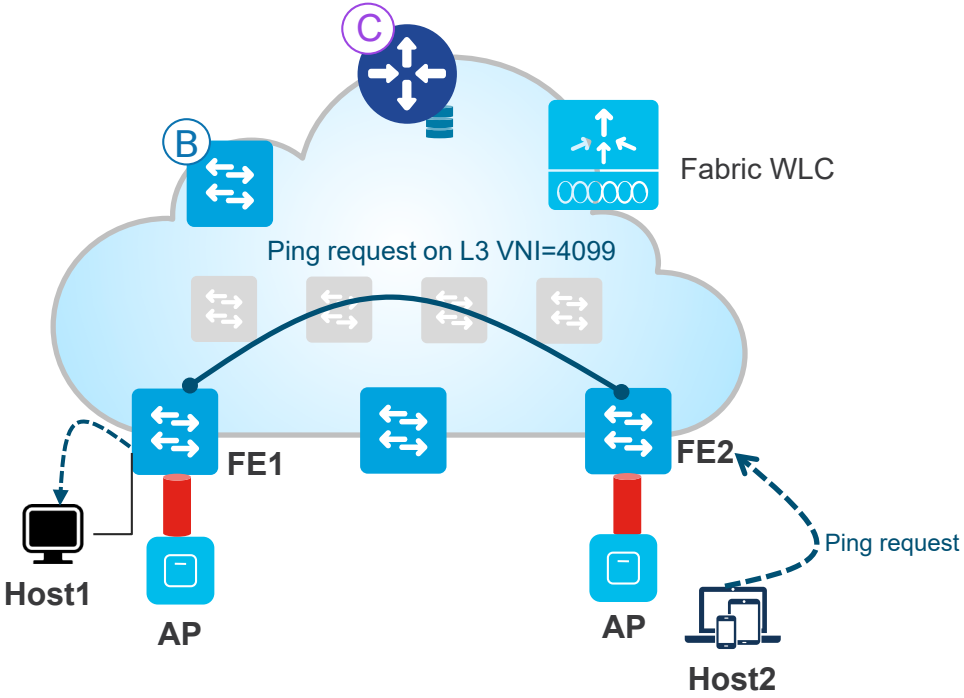

7 Once traffic reaches edge 1 it will decapsulate the VXLAN headers and based on the inner packet details will do a routing lookup to send it to Host1.

.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

1. Control Plane Lookup
2. Access Control Policy

Unicast Packet Forwarding & Access Control

Group-Based Policy

Ingress Classification & Egress Enforcement

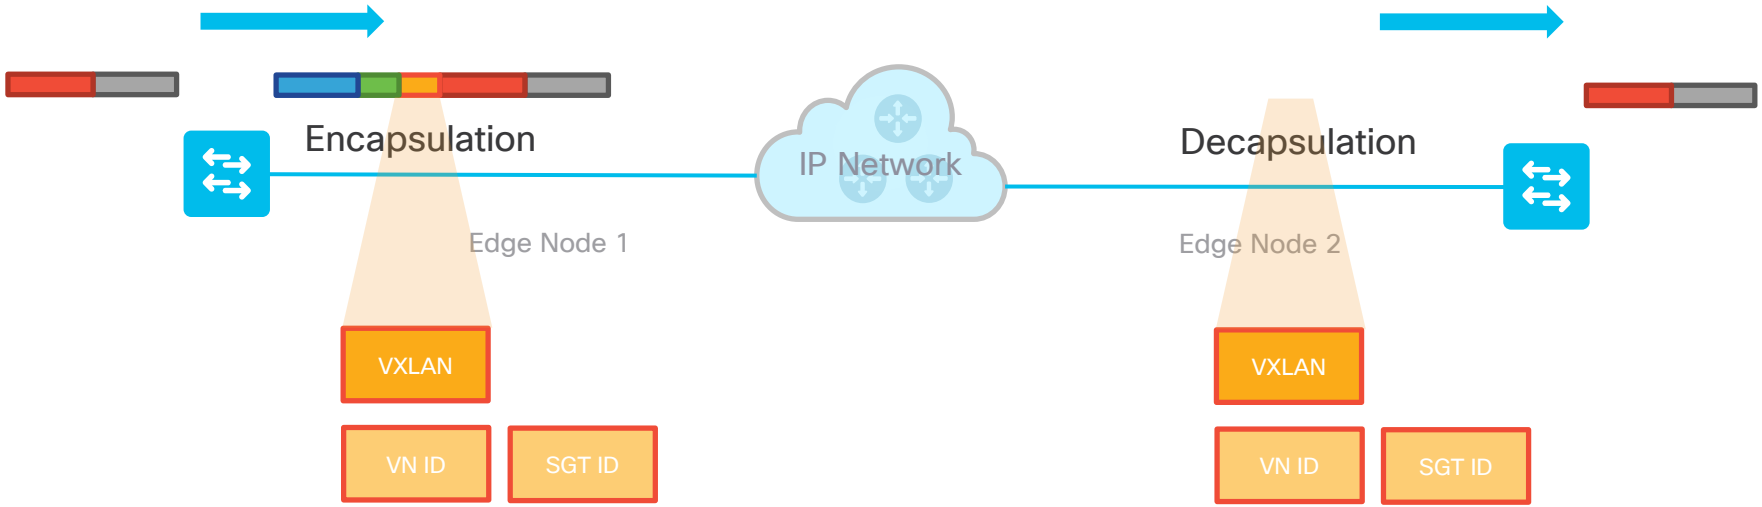

Classification
Static or Dynamic VN and SGT assignments

Propagation
Carry VN and Group context across the network

Enforcement
Group Based Policies ACLs, Firewall Rules

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy Definition in Cisco DNA Center

Policies in DNAC

REST API

SGACL matrix in ISE for rendering (runtime download)

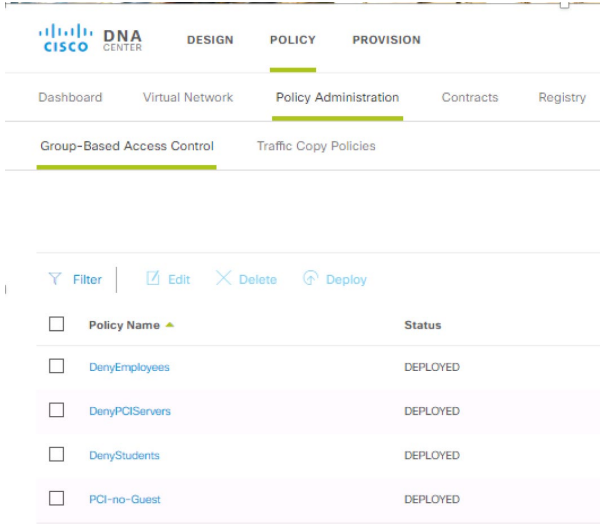

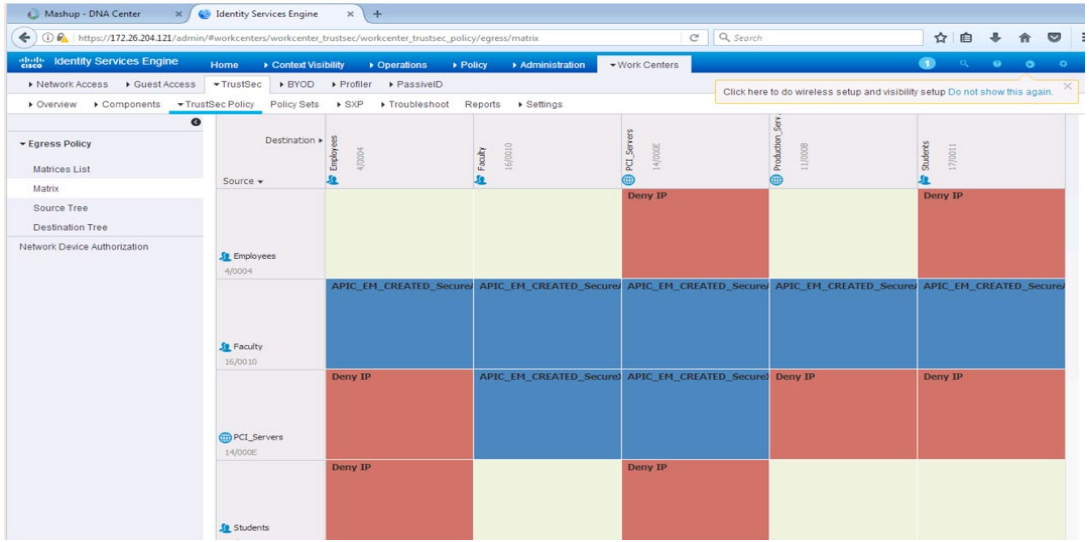

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Dynamic Policy Download

- New User/Device/Server provisioned
- Switch requests policies for assets they protect
- Policies downloaded & applied dynamically
- Result: All controls centrally managed
  - Security policies de-coupled from network topology
  - No switch-specific security configs needed
  - One place to audit network-wide policies

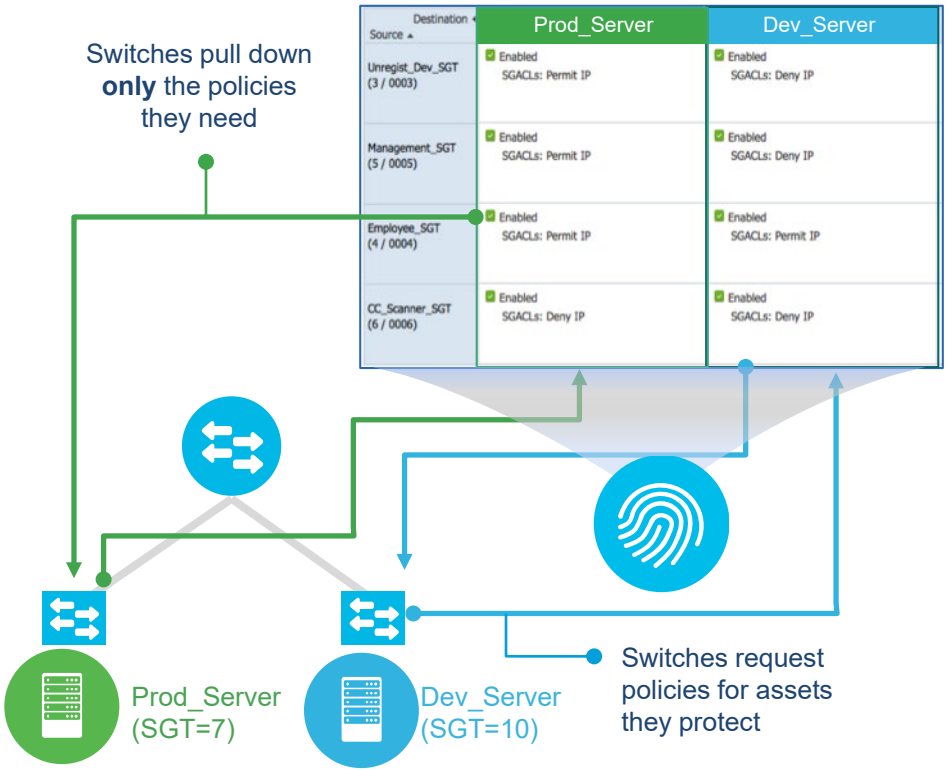

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Downloaded Policy

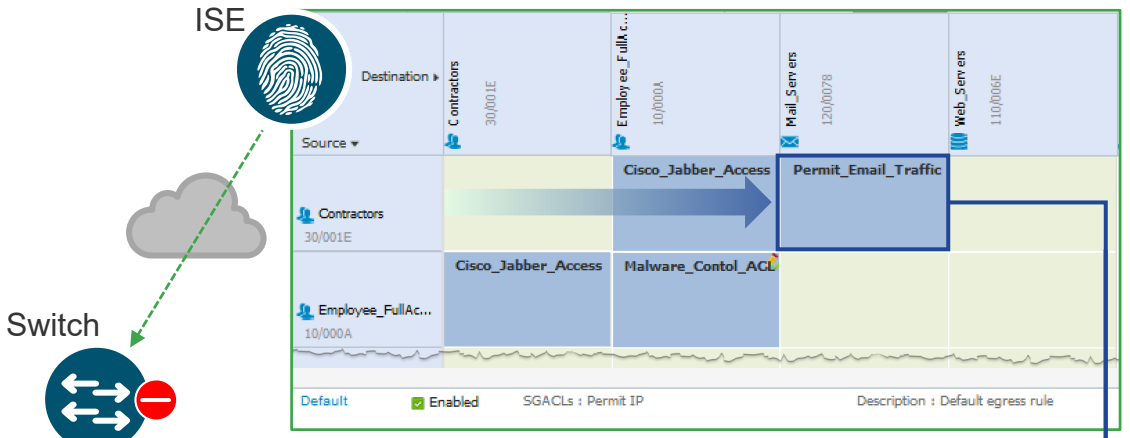


```
Switch#show cts rbacl Permit_Email_Traffic
CTS RBACL Policy
================
RBACL IP Version Supported: IPv4
  name    = Permit_Email_Traffic-40
  IP protocol version = IPV4
  refcnt = 1
  flag    = 0x40000000
  stale   = FALSE
  RBACL ACEs:
    permit tcp dst eq 110
    permit tcp dst eq 143
    permit tcp dst eq 25
    permit tcp dst eq 465
    permit tcp dst eq 585
    permit tcp dst eq 993
    permit tcp dst eq 995
    deny all log
```

```
Switch#show cts role-based permissions
IPv4 Role-based permissions default:
        Permit IP-00
...
IPv4 Role-based permissions from group 10:Employee_FullAccess to group
10:Employee_FullAccess:
        Malware_Contol_ACL-10
IPv4 Role-based permissions from group 10:Employee_FullAccess to group 30:Contractors:
        Cisco_Jabber_Access-10
IPv4 Role-based permissions from group 30:Contractors to group 10:Employee_FullAccess:
        Cisco_Jabber_Access-10
IPv4 Role-based permissions from group 30:Contractors to group 120:Mail_Servers:
        Permit_Email_Traffic
...
```

IOS switch as enforcer

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

1. Multicast modes
2. Broadcast forwarding

Advanced Topics

PIM ASM/SSM Control plane Interaction with the Fabric

Cisco SD-Access Fabric Architecture

Multicast with PIM-ASM – Multicast Receiver to RP

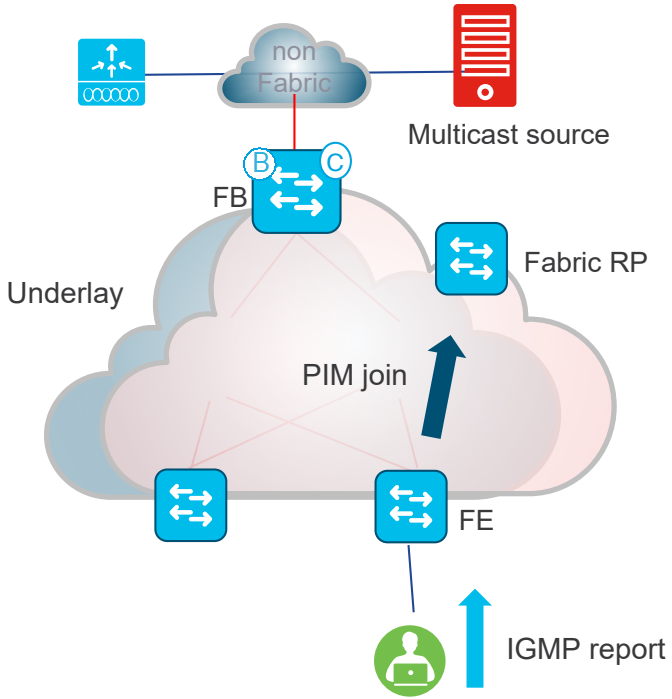

- Multicast client (receiver) is in the overlay, multicast source can be outside or inside the fabric
- PIM-ASM or SSM can be running in the overlay
- The client sends an IGMP report for a specific multicast group (G)
- The fabric Edge node (FE) sends a PIM join towards the Rendezvous Point RP
- The RP is registered as part of the end point IP (EID) space of the overlay.
- The edge node will ask the control plane for the location of the RP address (IP to RLOC Table) and based on the reply will send the PIM join in the overlay to the RP.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

How multicast works with PIM SSM

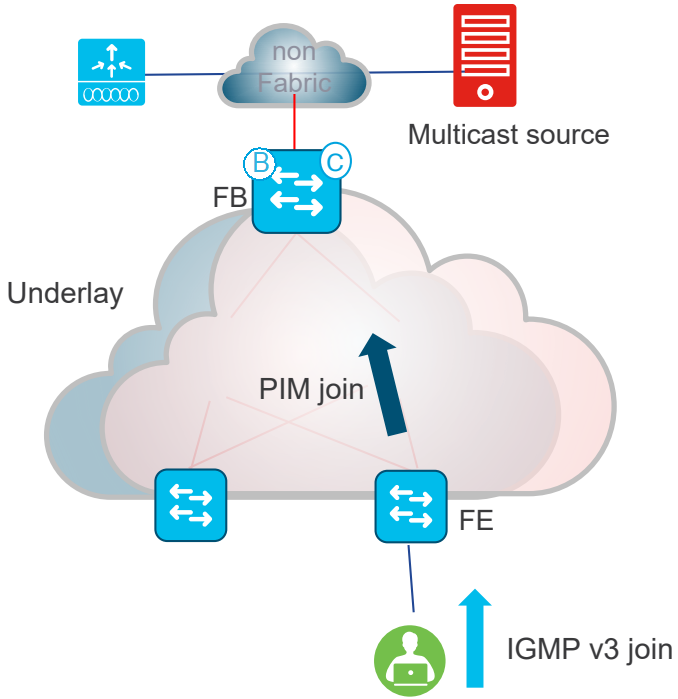

- Multicast client (receiver) is in the overlay, multicast source can be outside the fabric or in the overlay as well
- PIM-SSM needs to be running in the overlay
- An RP is not used in a PIM SSM deployment
- The client sends an IGMP v3 report for a specific multicast group (G)
- The fabric edge node (FE) receives it and since the IGMP v3 report has the source address information for that multicast group it sends a PIM join towards the source directly. In our case since the source is reachable through the border it sends the PIM join to the border.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

How multicast works with PIM SSM – Multicast Receiver to RP

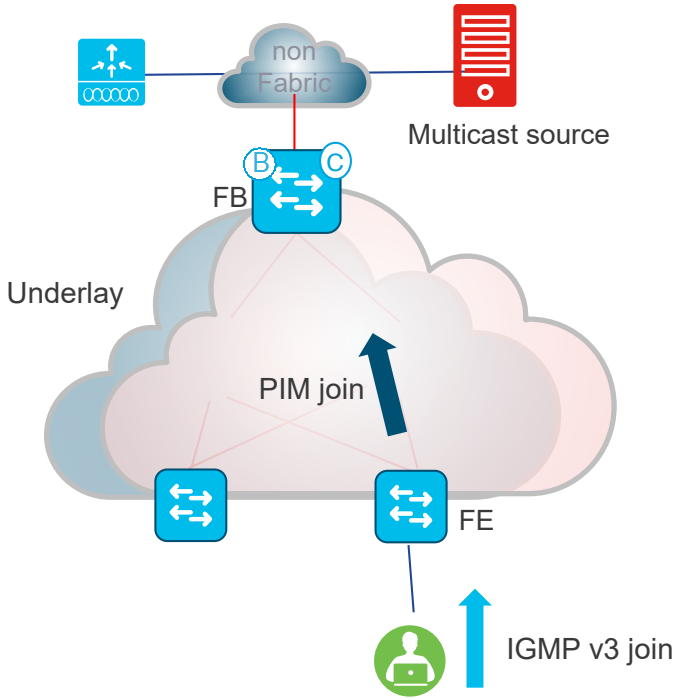

- Since in a SSM deployment the source address is part of IGMP v3 join the edge node will ask the control plane for the location of the source address (IP to RLOC Table) and based on the reply will send the PIM join in the overlay to the destination node.
- If Border registered that source then the PIM join is directly sent to the Border (the Border may forward the PIM join upstream towards the source, if the source is not directly connected).
- If the source is not known in the fabric, the PIM join is also sent to the border as it is the default exit point of the fabric.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

How multicast works with PIM ASM – Multicast Source to RP

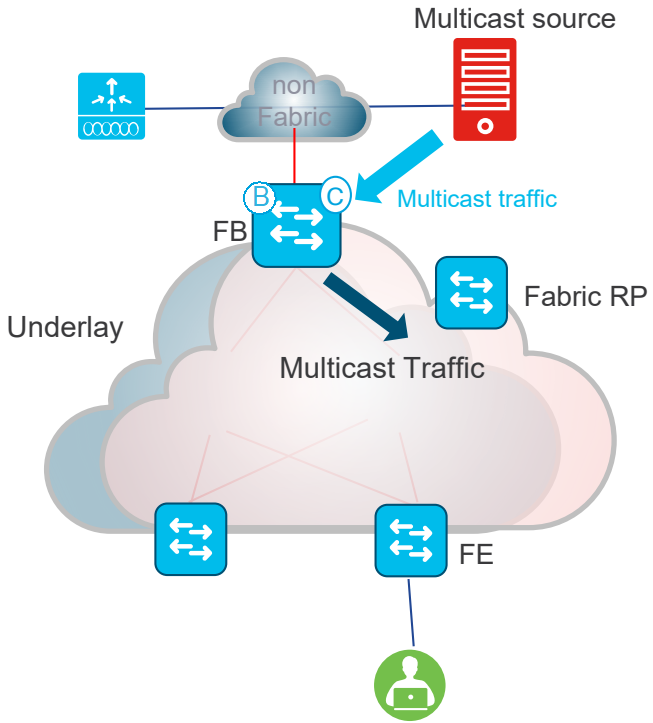

- The multicast source will send the multicast traffic on the interfaces towards the RP via the fabric Border(FB)
  - If the Source is directly connected to the border, the border is the DR for that segment
  - If the Source is not directly connected, the border is simply along the routed path towards the RP
- The FB receives it and sends the traffic towards the RP
- The Border node will ask the control plane for the location of the RP address (IP to RLOC Table) and based on the reply will send the traffic in the overlay to the RP
- The RP now has the source and receiver information for that multicast group

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

How multicast works with PIM SSM – Multicast Source to RP

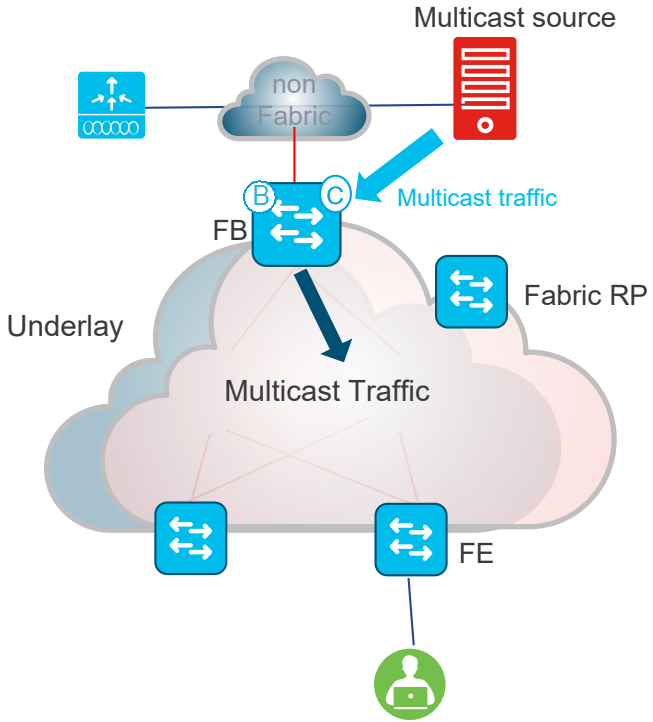

- The multicast source will send the multicast traffic on the interfaces towards the listeners via the fabric Border(FB)
  - If the Source is directly connected to the border, the border is the DR for the segment
  - If the Source is not directly connected, an SSM tree was formed when the PIM joins where forwarded previously and the Fabric Border is on the path of that tree.
- The FB receives it and sends the traffic towards the fabric Edge(s) as the PIM join is directly coming from the fabric Edge(s) to the Border in a SSM deployment.
- In a PIM SSM deployment there isn't an RP anchored shared tree .

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Head End Replication Multicast Data Plane in Fabric

Cisco SD-Access Fabric Architecture

How multicast works with PIM ASM – Data Plane

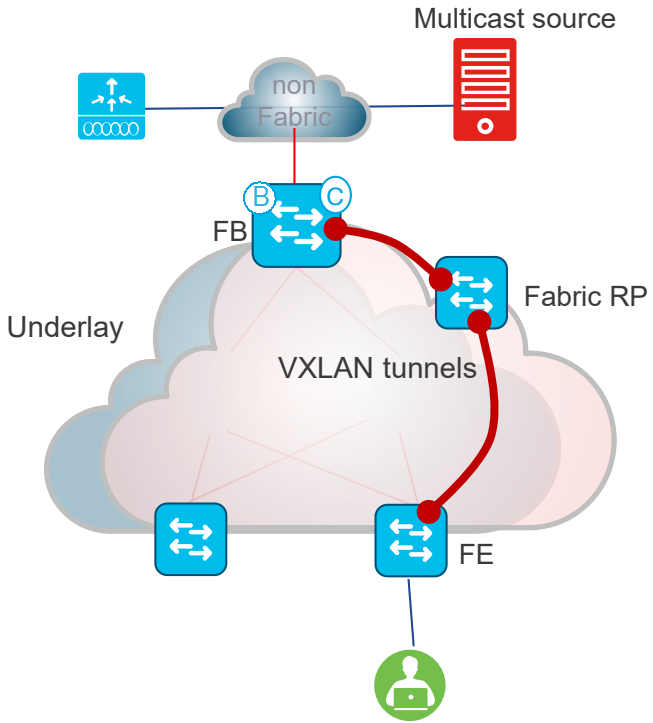

- The RP has the source and receiver information for a particular multicast group
- The fabric Border will send the multicast source traffic over a VXLAN tunnel to the RP. The RP will forward that traffic to FE over another VXLAN tunnel
- FE receives the VXLAN packets, decapsulates, applies policy and sends original IP multicast packet to the port on which the receiver is connected

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Multicast using PIM-ASM – Data Plane

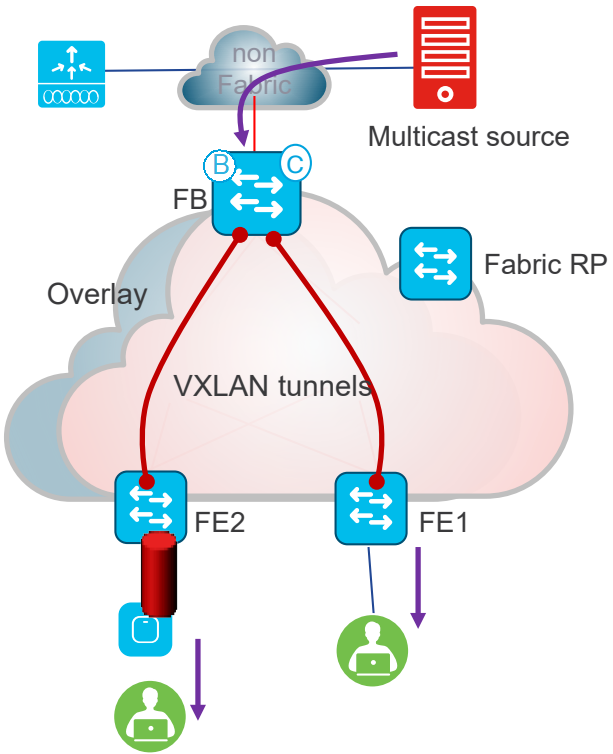

- Once the first multicast packet is delivered to the FE, the shortest path failover (SPT) takes place. Traffic is forwarded directly between the Border and the Edge.
- The FE knows that the Border owns the multicast source based on the first multicast packet received and send a PIM join directly to the Border for that multicast group.
- FB now knows which FEs have clients that requested the specific multicast group.
- It performs headend replication and VXLAN encapsulates the multicast traffic and unicasts it to the interested FEs
- The multicast traffic is sent in the overlay
- FE receives the VXLAN packets, decapsulates, policy and then sends original IP multicast packet to the port on which the receiver is connected.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Native Multicast

Cisco SD-Access Fabric Architecture

Native Multicast

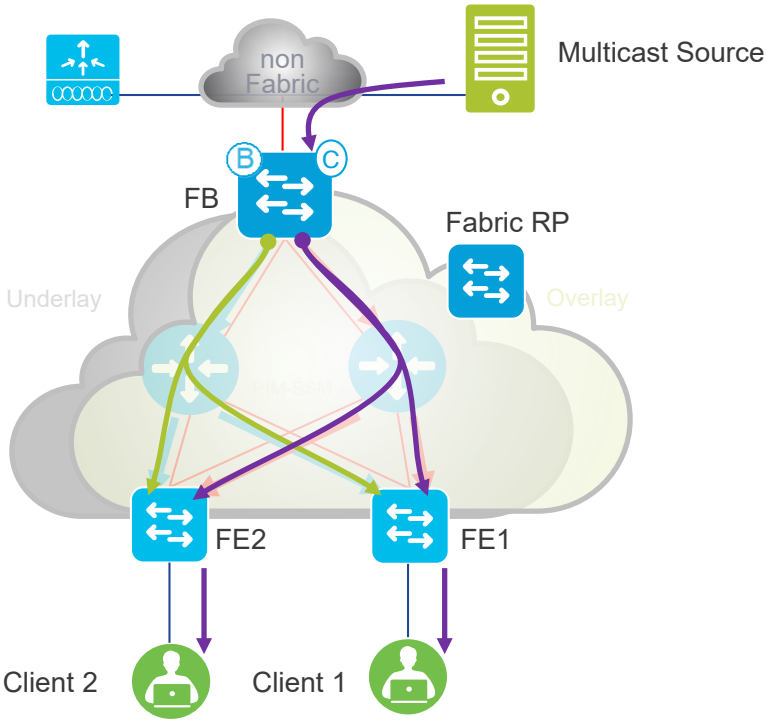

- Significantly reduces replication load at the Head-End
- Significantly improves overall scale and reduces latency

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

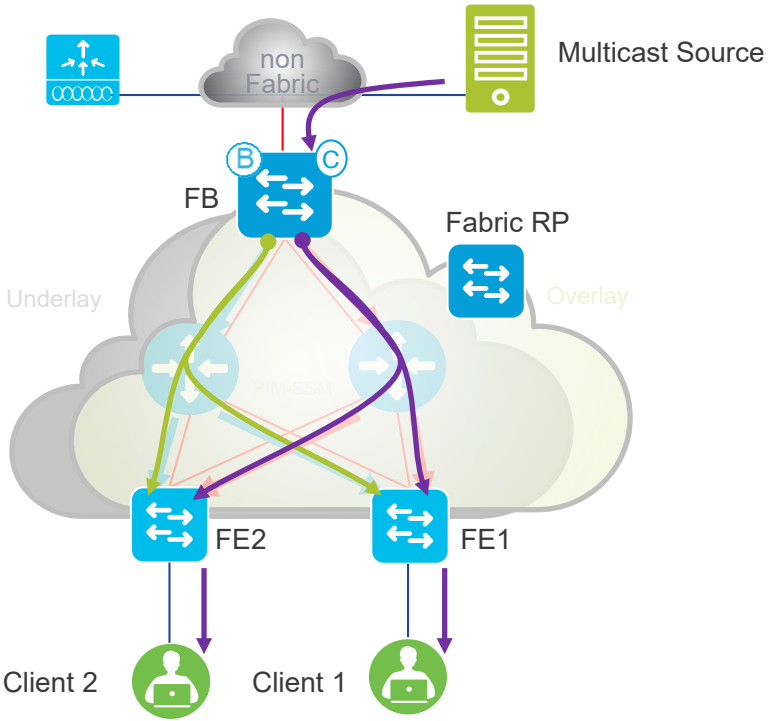

- All existing multicast control plane overlay behavior is the same
- PIM ASM and SSM can be used in the overlay as before.
- Each multicast group in the overlay is mapped to a corresponding (PIM SSM) multicast underlay group
- Multicast distribution (replication) occurs natively within the underlay network (e.g. intermediate nodes)
- Incoming multicast traffic for a given VN is encapsulated in VXLAN, and then sent with {Source IP = FE node RLOC, Destination IP = Underlay Multicast Group} as the outer IP addresses.
- PIM SSM is used in the underlay for multicast transport

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

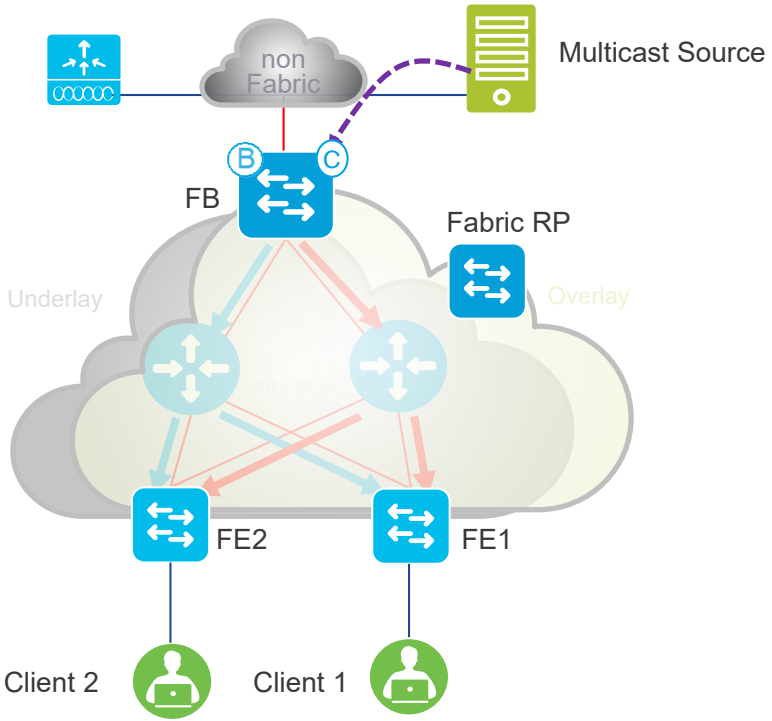

When the native multicast knob is turned on for a given fabric site the VNs where multicast is turned on will be instructed to move over to native multicast for the data path.

The configuration is pushed under the LISP interface for the respective VNs and the multicast groups in that VN will be mapped to underlay SSM groups for data transport.

Interface LISP0.4096
 ip pim lisp transport multicast
 ip pim lisp core-group-range 232.0.0.1 1000

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

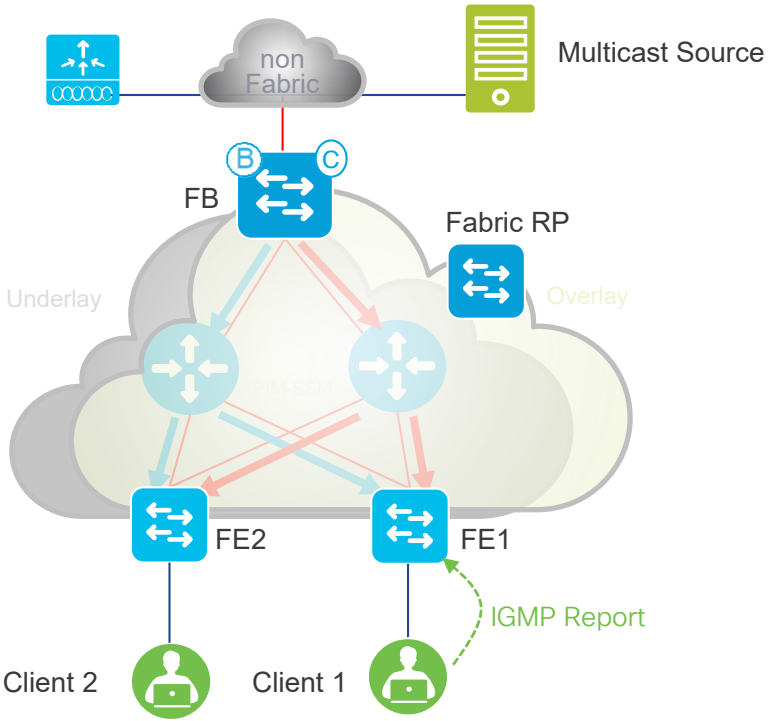

1. Client 1, a multicast receiver sends an IGMP report for group 238.0.0.1 to the fabric Edge

In this example we assume that ASM is used in the overlay and the group address is 238.0.0.1.

The overlay ASM group is mapped to SSM group (RP-RLOC, 232.0.0.9) in the underlay.

This is derived based on the configuration

Since we are using SSM in the underlay for native multicast there is no pre built multicast tree for any given group in overlay.

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

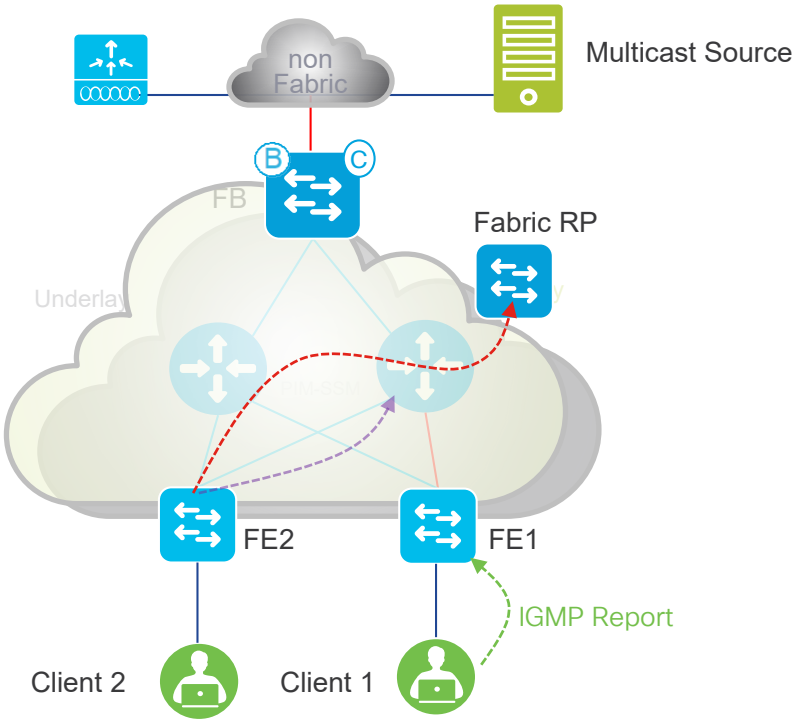

2 The fabric Edge node receives the IGMP report for group 238.0.0.1, it sends the corresponding PIM joins:

- in the overlay: a PIM join to the RP for 238.0.0.1
- In the underlay: a PIM join for (RP-RLOC,232.0.0.9) (SSM group).

The source address in the underlay SSM join will be the RLOC address of the RP. The Information of the RLOC address for the RP is provided by the control plane node.

Native multicast maps an overlay ASM group to an SSM underlay multicast group.

Thus, the *,G joins in the overlay and RP-RLOC,G joins in the underlay are assembled and transmitted.

A multicast tree anchored at the RP-RLOC is created in the underlay

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

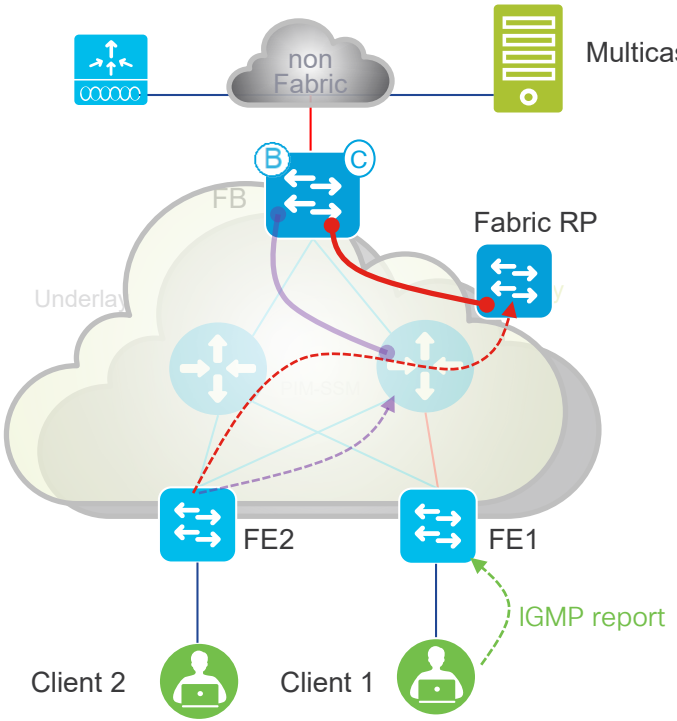

3

The multicast source starts sending traffic.

The overlay multicast traffic is unicast encapsulated to the RP-RLOC, it is then multicast encapsulated in the (S,G) tree from the RP-RLOC onwards.

The iTR (in our case the fabric Border) will send a source registration message in the overlay on the group address 238.0.0.1 to the RP and also sends the 238.0.0.1 multicast traffic in the overlay to the RP-RLOC (unicast encapsulated).

The RP-RLOC xTR sends the traffic in the underlay on the mapped group 232.0.0.9.

The information of the RLOC address is provided by control plane node. The control plane gives the IP to RLOC information for the RP address.

Up to this point an RP anchored ASM tree has been created

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-Access Fabric Architecture

Native Multicast

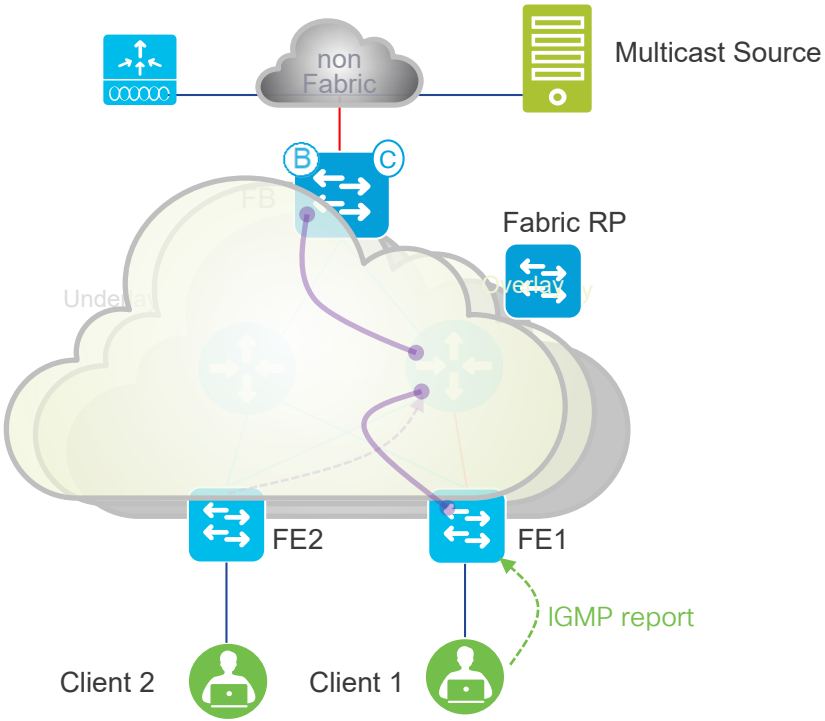

4 The underlay now has enough information to replicate the traffic to the needed devices.

When SPT failover takes place, the source specific PIM joins in the overlay are mapped to source specific PIM joins in the underlay and an SSM SPT to the actual source is formed in the underlay.

Pre-SPT: Underlay (RP-RLOC, 232.0.0.9)
Post SPT: Underlay (Source-RLOC, 232.0.0.9)

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Summary

SD-Access Support

Digital Platforms for your Cisco Digital Network Architecture

For more details: cs.co/sda-compatibility-matrix

Switching

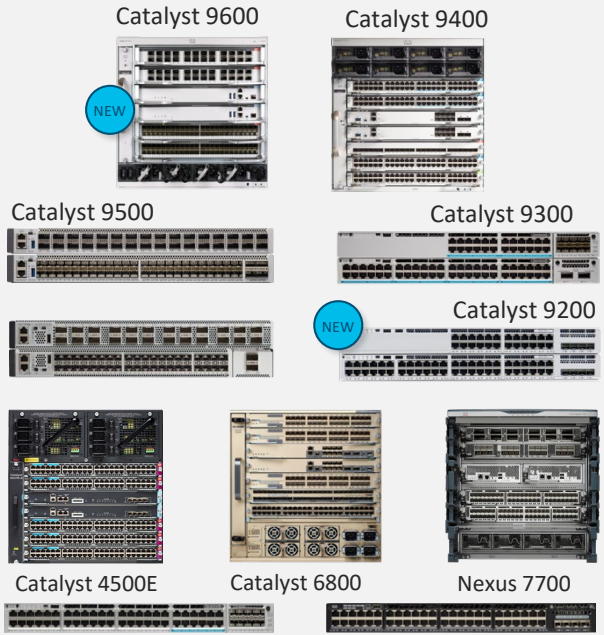

Routing

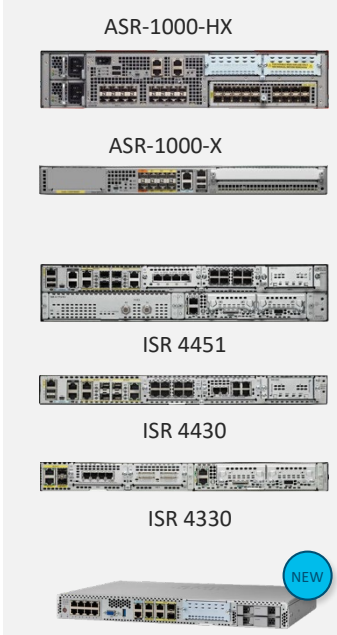

Wireless

Extended BETA

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

What to Do Next?

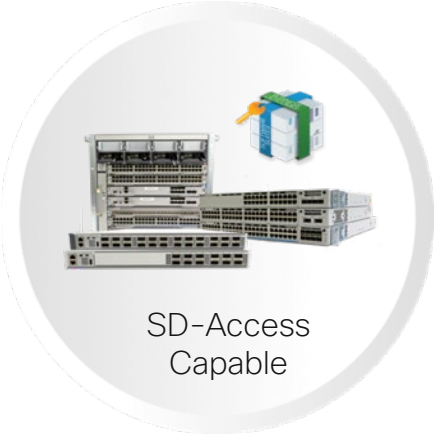

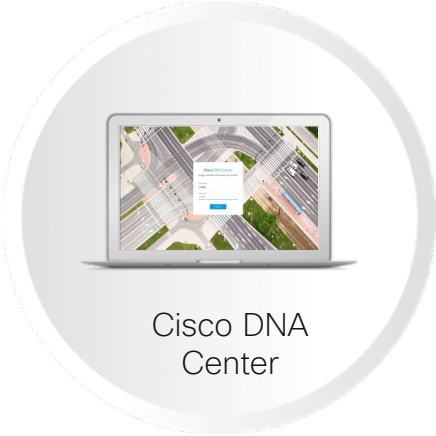

Refresh your **Hardware & Software**	Deploy the **Cisco DNA Center**	Engage with **Cisco Services**
Get **SD-Access supported Devices** with **Cisco DNA Advantage Licenses**	Get **Cisco DNA Center Appliances** with **Cisco DNA Center Software**	**Cisco Services** can help you **Test - Migrate - Deploy - Manage**

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-Access Resources

Would you like to know more?

cisco.com/go/dna

cisco.com/go/sdaccess

- SD-Access At-A-Glance
- SD-Access Ordering Guide
- SD-Access Solution Data Sheet
- SD-Access Solution White Paper

cisco.com/go/cvd

- SD-Access Design Guide
- SD-Access Deployment Guide
- SD-Access Segmentation Guide

cisco.com/go/dnacenter

- Cisco DNA Center At-A-Glance
- Cisco DNA ROI Calculator
- Cisco DNA Center Data Sheet
- Cisco DNA Center 'How To' Video Resources

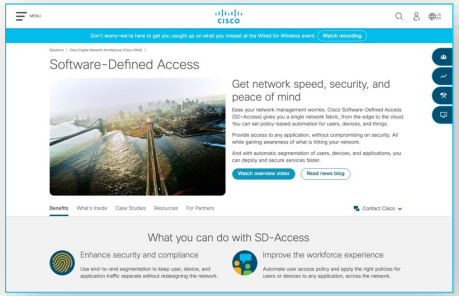

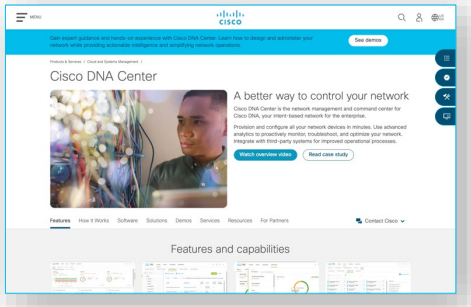

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

More Resources

Understand the underpinnings of your SDA fabric

Understand the potential of your LISP enabled network

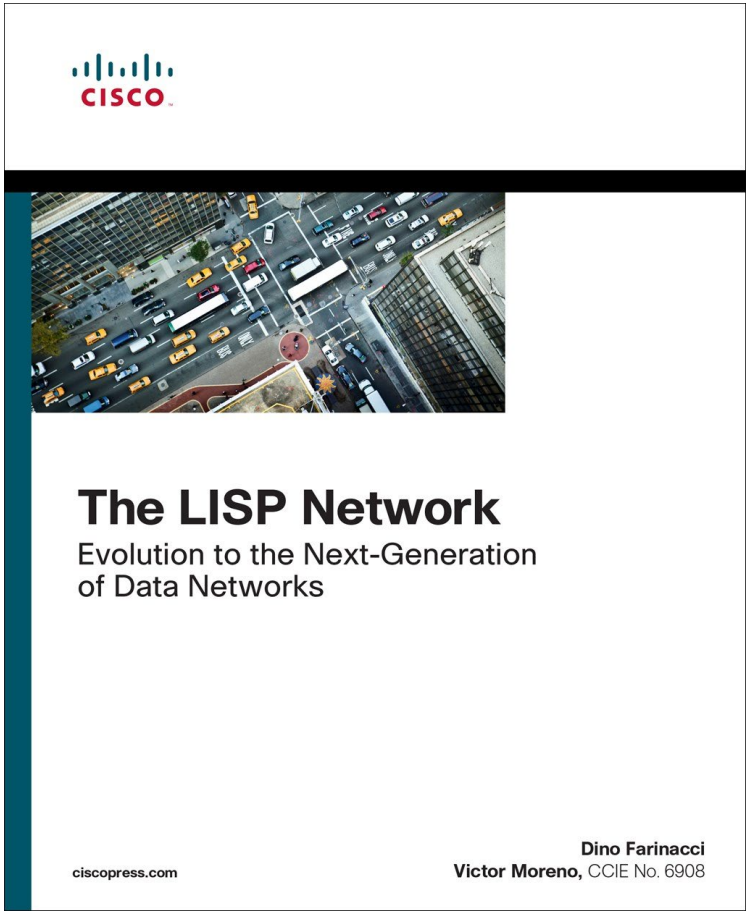

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

cisco Live!

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Continue your education

Demos in the Cisco Showcase

Walk-In Labs

Meet the Engineer 1:1 meetings

Related sessions

© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public

Thank you

You make **possible**