cisco Live!

Let's go

The bridge to possible

Design, Automate, and Manage VXLAN BGP EVPN Multi–Site with NDFC

Parth Patel, Technical Leader, Technical Marketing Engineer:
Data Center and Provider Connectivity BU

BRKDCN-2128

Cisco Webex App

Questions?

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

How

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

Webex spaces will be moderated by the speaker until December 22, 2023.

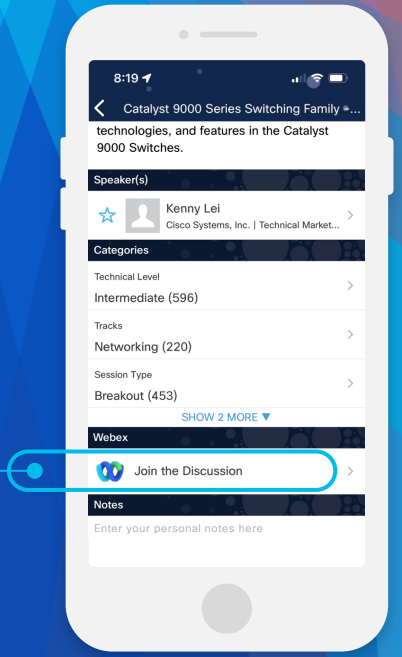

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2128

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- VXLAN EVPN Multi-Site Overview
- Design VXLAN EVPN Multi-Site with NDFC
- Why NDFC?
- Automate and Manage VXLAN EVPN Multi-Site with NDFC
- Conclusion

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN EVPN Multi-Site Architecture

VXLAN EVPN Multi-Site Overview

Functional Components

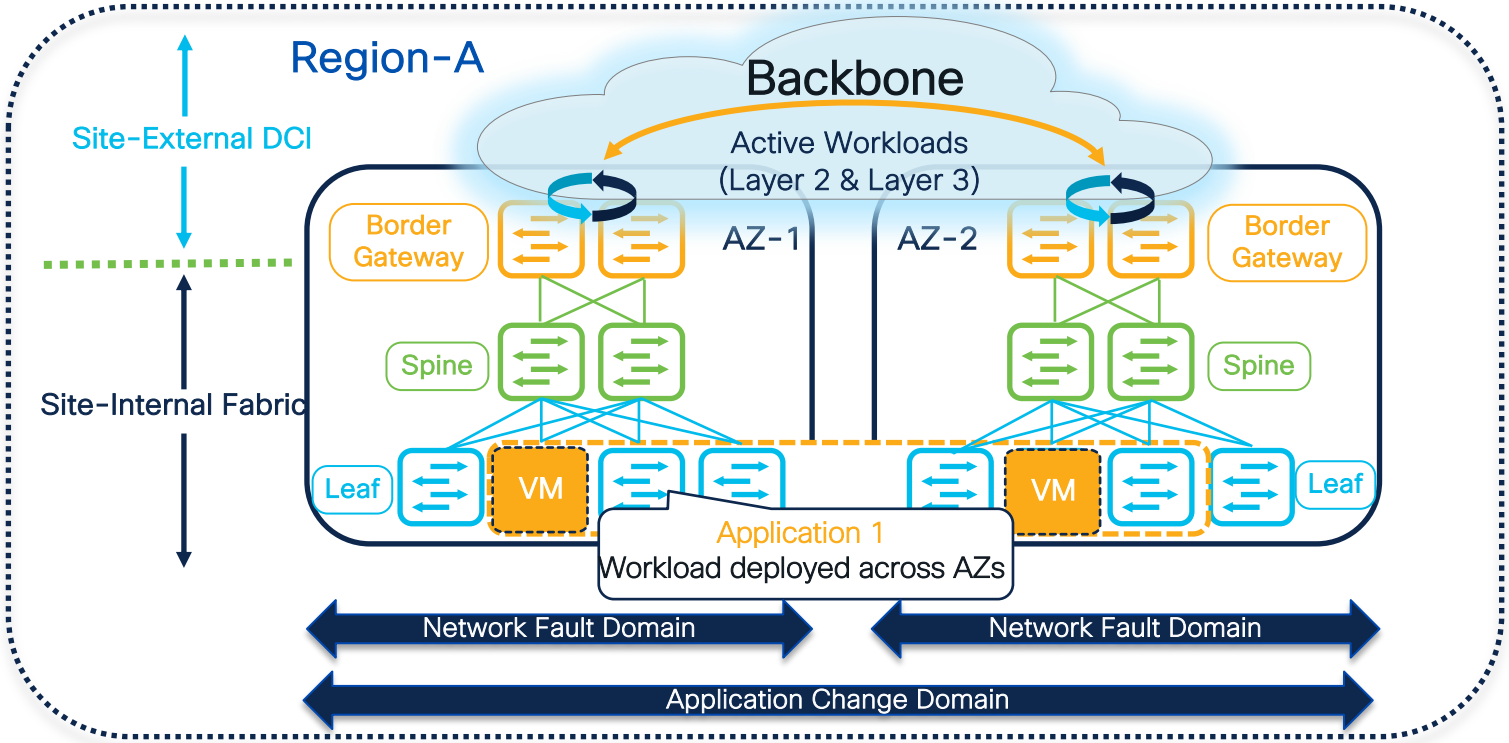

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW to Cloud

- Scalable design option with horizontal scale-out within and across multiple sites.
- Backbone/Cloud can be any routed service. (flat L3, MPLS-L3VPN)
- **Multi-Site Underlay: eBGP IPv4 Unicast** ⟷
  - Site-External DCI BUM: Ingress-Replication or Multicast supported. *Currently NDFC supports only Ingress-Replication.
  - Site-Internal Fabric BUM: Ingress-Replication or Multicast supported independently at each site.
- **Multi-Site Overlay: eBGP EVPN Overlay** ⟷
  - Full-mesh BGP EVPN peering across all BGWs.
- Ensure that PIP / VIP of all BGWs are known by every BGW and MTU must accommodate VXLAN encapsulated traffic.

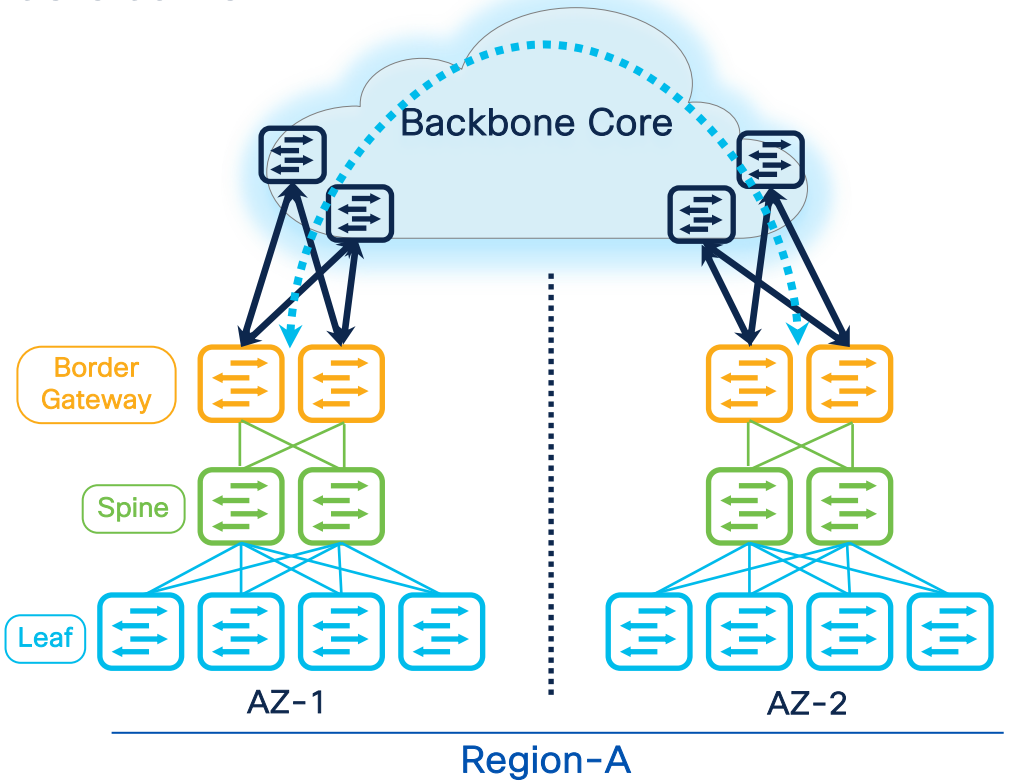

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW Back-to-Back

- Available option for connecting 2 sites.
- Minimum topology is the square. An enhanced option is to add links between site local BGWs for improved ECMP and Failure scenarios.
- **Multi-Site Underlay: eBGP IPv4 Unicast** ←——→
  - Site-External DCI BUM: Ingress-Replication or Multicast supported. *Currently NDFC supports only Ingress-Replication.
  - Site-Internal Fabric BUM: Ingress-Replication or Multicast supported independently at each site.
- **Multi-Site Overlay: eBGP EVPN Overlay** ←·····→
  - Full-mesh BGP EVPN peering across all BGWs.
- Ensure that PIP / VIP of all BGWs are known by every BGW and MTU must accommodate VXLAN encapsulated traffic.

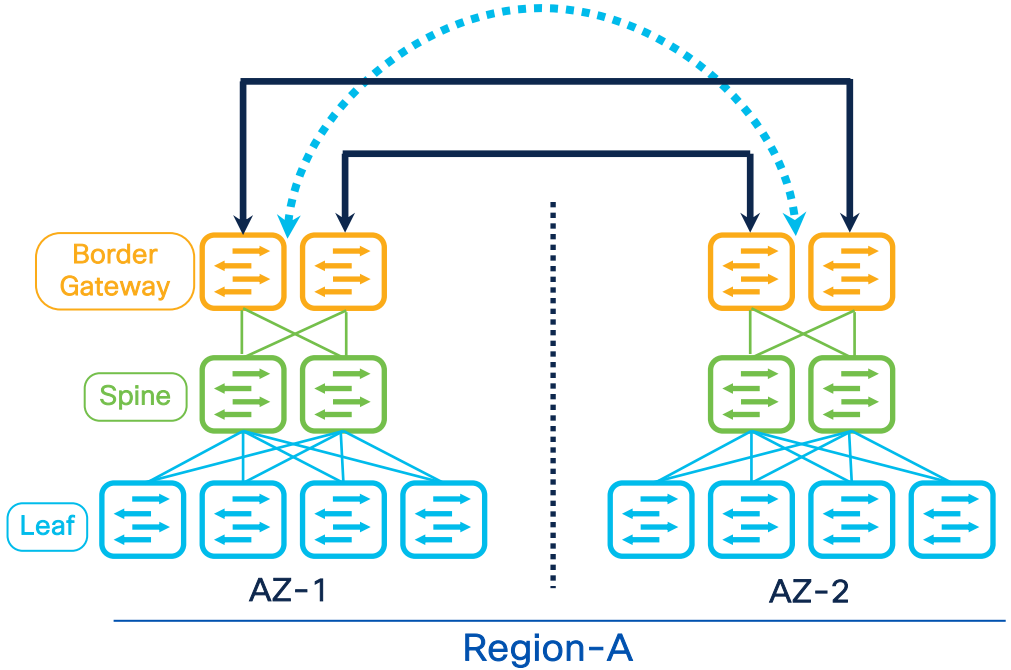

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multi-Destination Traffic Forwarding

EVPN Route-Type 4 (Ethernet Segment Route)

- BGW advertises BGP EVPN Type-4 update to all of their BGP EVPN peers (Site-Internal and Site-External)
- BGW will use this information for DF (Designated Forwarded) election on per VLAN/VNI
- EVPN Type-4 update will also carry ES Import Route-Target Extended Community
- Remote site BGWs will not install the route due to different RT and Site ID
- Command to check: **`show bgp l2vpn evpn route-type 4`**
  - BGP routing table entry for [4]:[0300.0000.00fd.e900.0309]:[32]:[10.1.1.1]/136
  - [4] = EVPN Route-Type 4
  - [03]= ESI Type MAC-based
  - [0000.0000.fde9] = ESI system MAC for Site-ID and RT {fde9 = 65001}
  - [000309] = ESI local discriminator
  - [32] = IP Address Length
  - [10.1.1.1] = Sender (BGW Primary IP) used for DF election process

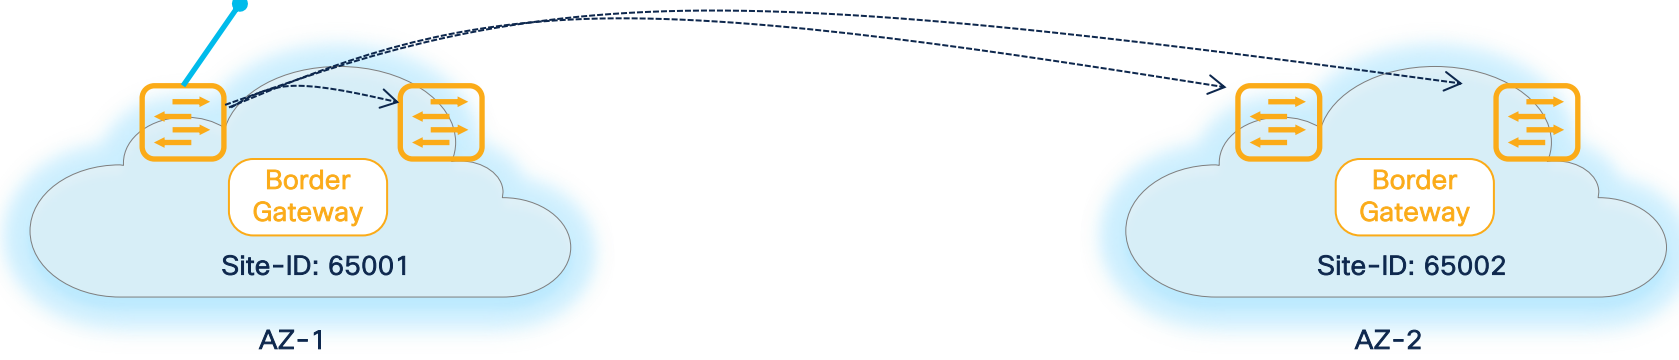

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multi-Destination Traffic Forwarding

BGW Designated Forwarded (DF) Election

- BGW learns about other BGWs originator IPs on the same site from exchanging EVPN Type 4 routes. The originator IP address used is the BGW NVE source interface's primary IP address.
- Each BGW creates an ordinal list of originator IP in numerical order from lowest to highest. Every BGW is then given an ordinal value based on its position in the ordinal list starting from position 0. The BGW with the lowest originator IP would get an ordinal value of 0. The ordinal value decides which BGW will be the DF for a VLAN/VNI.
- DF election Formula uses “mod” operation:
- `I = V mod N {I = ordinal value, V = VLAN #, N = # BGWs in a site}`
- If there are 4 BGWs in one site. The ordinal list of the BGWs will be arranged as shown below.
- Ordinal List = `10.1.1.1, 10.1.1.2, 10.1.1.3, 10.1.1.4 {I=0, I=1, I=2, I=3}`
- The DF for VLAN 48 will be: “`I = 48 mod 4 = 0 = 10.1.1.1`”
- The DF for VLAN 49 will be: “`I = 49 mod 4 = 1 = 10.1.1.2`”
- The DF for VLAN 50 will be: “`I = 50 mod 4 = 2 = 10.1.1.3`”
- The DF for VLAN 51 will be: “`I = 51 mod 4 = 3 = 10.1.1.4`”

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Multi-Destination Traffic Forwarding

EVPN Route-Type 3 (Inclusive Multicast Ethernet Tag Route)

- BGW advertises BGP EVPN Type-3 only towards Site-External DCI
- Receiving BGW switch does not forward it to Site-Internal Fabric
- Command to check: **`show bgp l2vpn evpn route-type 3,`**
- **`show l2route evpn imet all detail`**
- **`Local: Advertising from Site 1 BGW to Site 2 BGW`**

```
BGP routing table entry for [3]:[0]:[32]:[10.3.0.4]/88
AS-Path: NONE, path locally originated
Extcommunity: RT:65001:10000 ENCAP:8
  PMSI Tunnel Attribute:
    flags: 0x00, Tunnel type: Ingress Replication
    Label: 10000, Tunnel Id: 10.3.0.4
Path-id 1 advertised to peers: 20.2.0.4
```

- **`Remote: Learning from Site 2 BGW to Site 1 BGW`**

```
BGP routing table entry for [3]:[0]:[32]:[20.3.0.4]/88
Path type: external, path is valid, is best path
 Imported to 1 destination(s)
 Imported paths list: L2-10000
AS-Path: 65002 , path sourced external to AS
Extcommunity: RT:65001:10000 ENCAP:8
  PMSI Tunnel Attribute:
    flags: 0x00, Tunnel type: Ingress Replication
    Label: 10000, Tunnel Id: 20.3.0.4
```

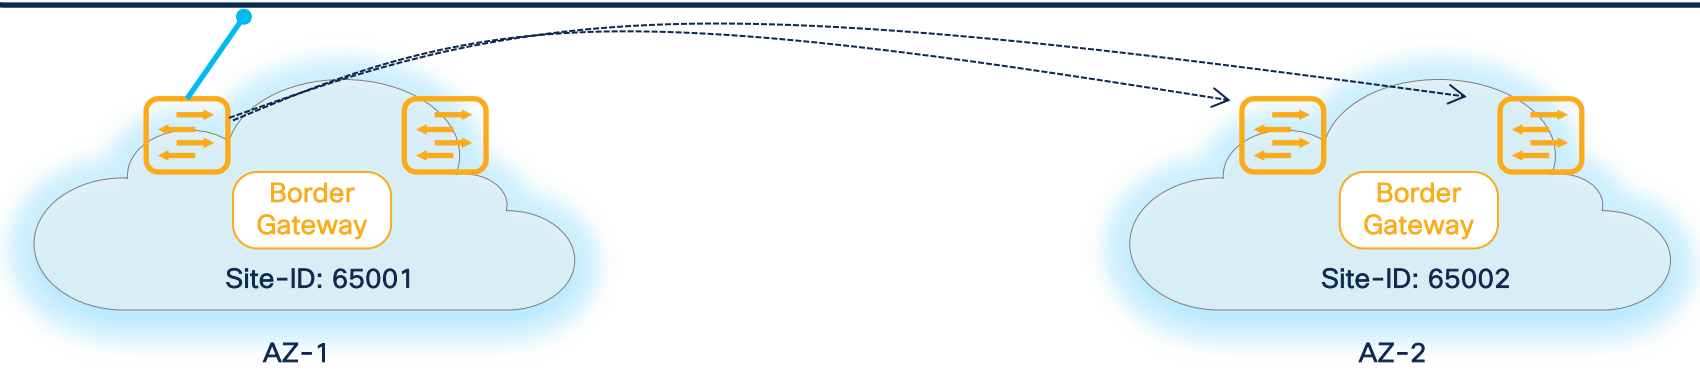

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW Back-to-Back BUM (w/o Local L3 link)

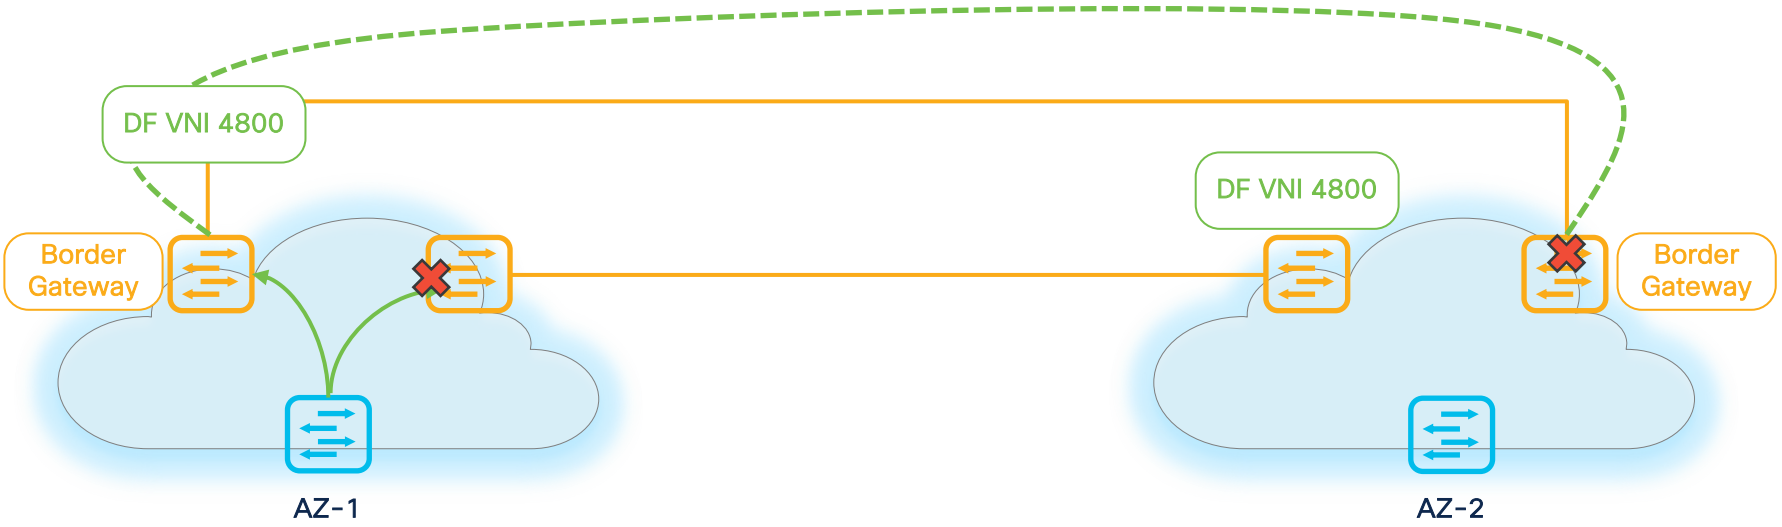

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW Back-to-Back BUM (with Local L3 link)

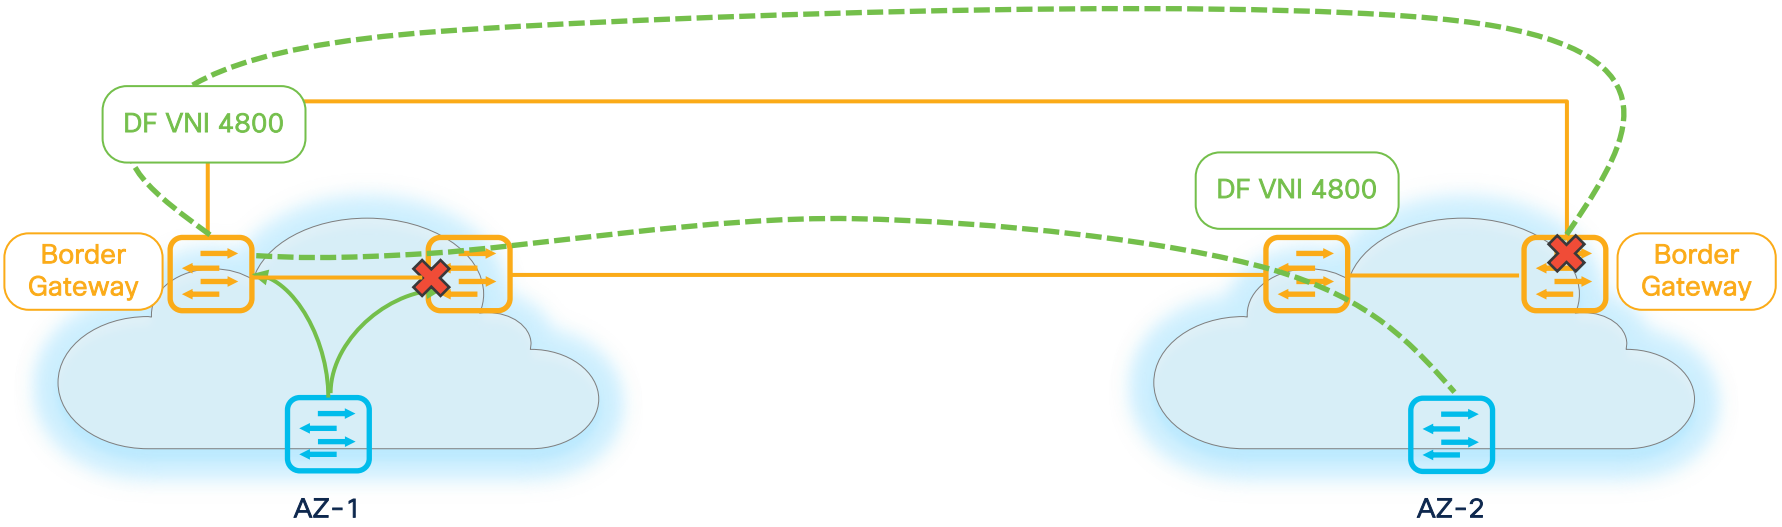

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW to Centralized Router-Server

- BGP Route-Server (RS) function can be used in a highly scalable design to act as the central EVPN peering point for the Multi-Site.
- RS does not need to be on the data path.
- **Multi-Site Underlay: eBGP IPv4 Unicast** ⟷
  - Site-External DCI BUM: Ingress-Replication or Multicast supported. *Currently NDFC supports only Ingress-Replication.
  - Site-Internal Fabric BUM: Ingress-Replication or Multicast supported independently at each site.
- **Multi-Site Overlay: eBGP EVPN Overlay** ⟷
  - BGP EVPN peering only between BGWs and RS.
- **Route-Server (RS) must support...**
  - EVPN AFI and Router-Server function per RFC 7947
  - Next-hop-unchanged, Retain RT, and RT Rewrite function
- Ensure that PIP / VIP of all BGWs are known by every BGW and MTU must accommodate VXLAN encapsulated traffic.

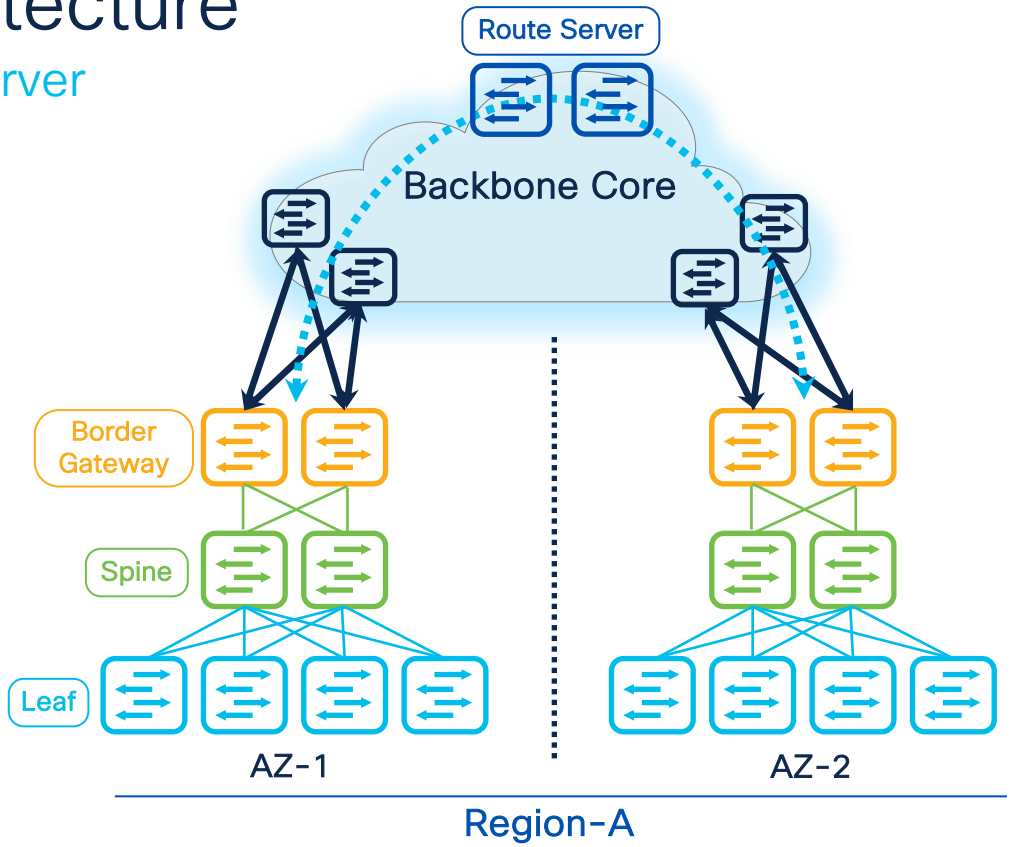

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN Multi-Site Architecture

DCI- BGW to Centralized Router-Server

- BGP Route-Server (RS) function can be used in a highly scalable design to act as the central EVPN peering point for the Multi-Site.
- RS can also be on the data path.
- **Multi-Site Underlay: eBGP IPv4 Unicast** ⟷
  - Site-External DCI BUM: Ingress-Replication or Multicast supported. *Currently NDFC supports only Ingress-Replication.
  - Site-Internal Fabric BUM: Ingress-Replication or Multicast supported independently at each site.
- **Multi-Site Overlay: eBGP EVPN Overlay** ⟷
  - BGP EVPN peering only between BGWs and RS.
- **Route-Server (RS) must support…**
  - EVPN AFI and Router-Server function per RFC 7947
  - Next-hop-unchanged, Retain RT, and RT Rewrite function
- Ensure that PIP / VIP of all BGWs are known by every BGW and MTU must accommodate VXLAN encapsulated traffic.

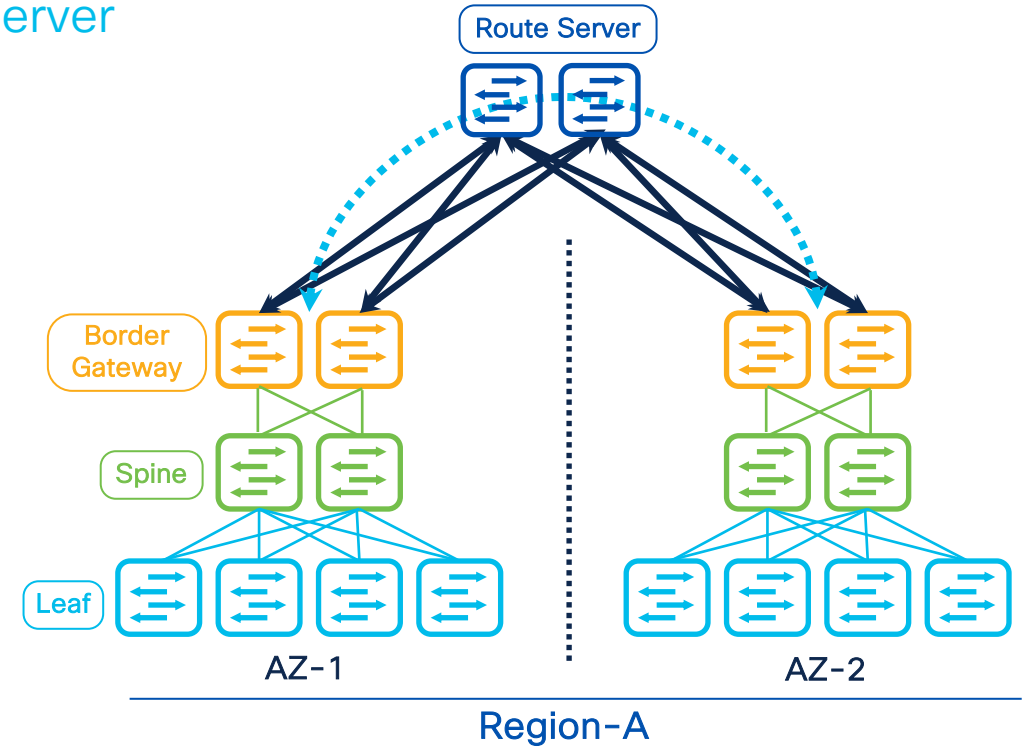

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cool! But What's the Catch?

There is always a catch ☺

1. This seems really cool! Is it easy to configure?
   There are lots of moving parts: OSPF/BGP/VXLAN. Manual configuration can be challenging.

2. How easy is it to make changes?
   You still rely on traditional SSH based management to each device, which can be cumbersome and error prone.

3. How much Visibility do I have into the network?
   Visibility and Troubleshooting is still performed on a "switch-by-switch" basis.

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Why do YOU need NDFC?

Why NDFC?

Multi-Architecture

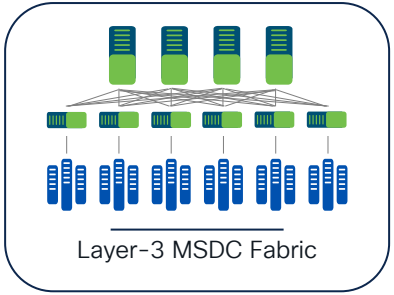

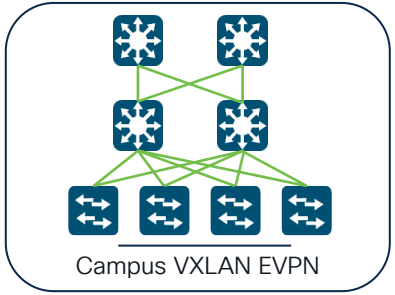

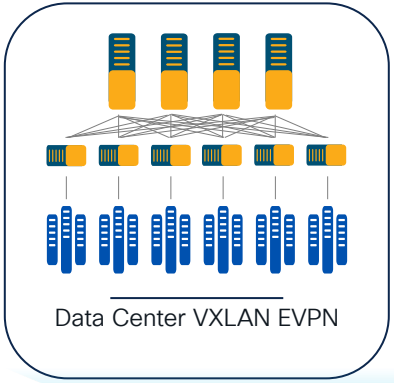

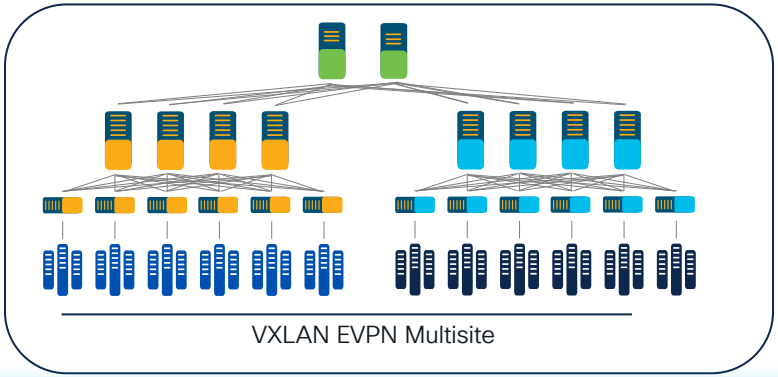

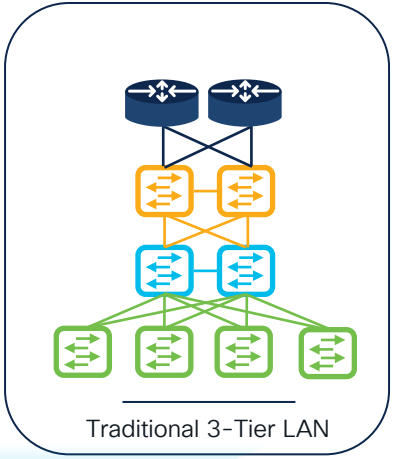

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Why NDFC?

Multi-Topology, Multi-Protocol

Rich set of control plane and data plane possibilities available

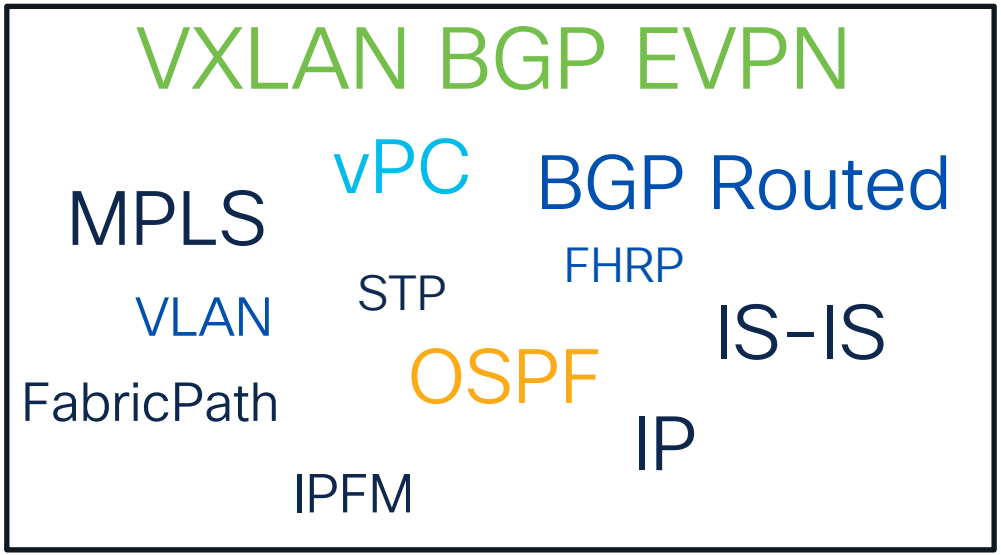

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Why NDFC?

Multi-Domain, Multi-Platform

NX-OS Nexus 9000 and 3000

IOS-XE Catalyst 9000

IOS-XR ASR 9000

NX-OS Nexus 7000

IOS-XE ASR 1000

TIP

Supported Hardware and Software might vary depending on NDFC version

Check compatibility matrix 12.1.3b

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Manage and Deploy VXLAN EVPN Multi-Site with NDFC

cisco Live!

VXLAN EVPN Multi-Site deployment with NDFC

Deploy in a few minutes!

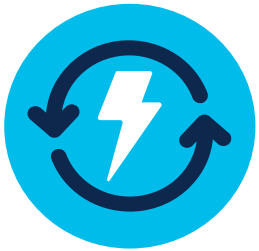

1. Deploy your individual DC fabrics (Data Center VXLAN EVPN)
2. (Optional) Deploy a fabric for Route-Server (Multi-Site Interconnect Network)
3. Deploy a fabric for the Multi-Site (VXLAN EVPN Multi-Site) & import all the previous fabrics into this one
4. Recalculate and deploy in the VXLAN EVPN Multi-Site Fabric

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN BGP EVPN Greenfield

Not on VXLAN EVPN Today?

- NDFC Fabric Controller Mode
- Build VXLAN fabric in few minutes
- Templates already embed best practices
- IP addresses, overlay pool, routing profiles, replication attributes –all taken care by NDFC

Step 1
Create

Define fabric settings (Underlay, Overlay) – AS#, Replication Mode, IGP, IP Pools, etc.

Step 2
Discover

Import switches with POAP or Day-0 config

Define switch Roles (Border, Leaf, Spine, etc)

[Optional] Create vPC pairs

Step 3
Recalculate and Deploy

Generates config based on intent

Preview side by side diffs

Cisco best practice implemented

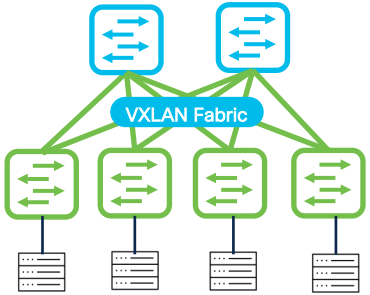

Fast, automated process

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric

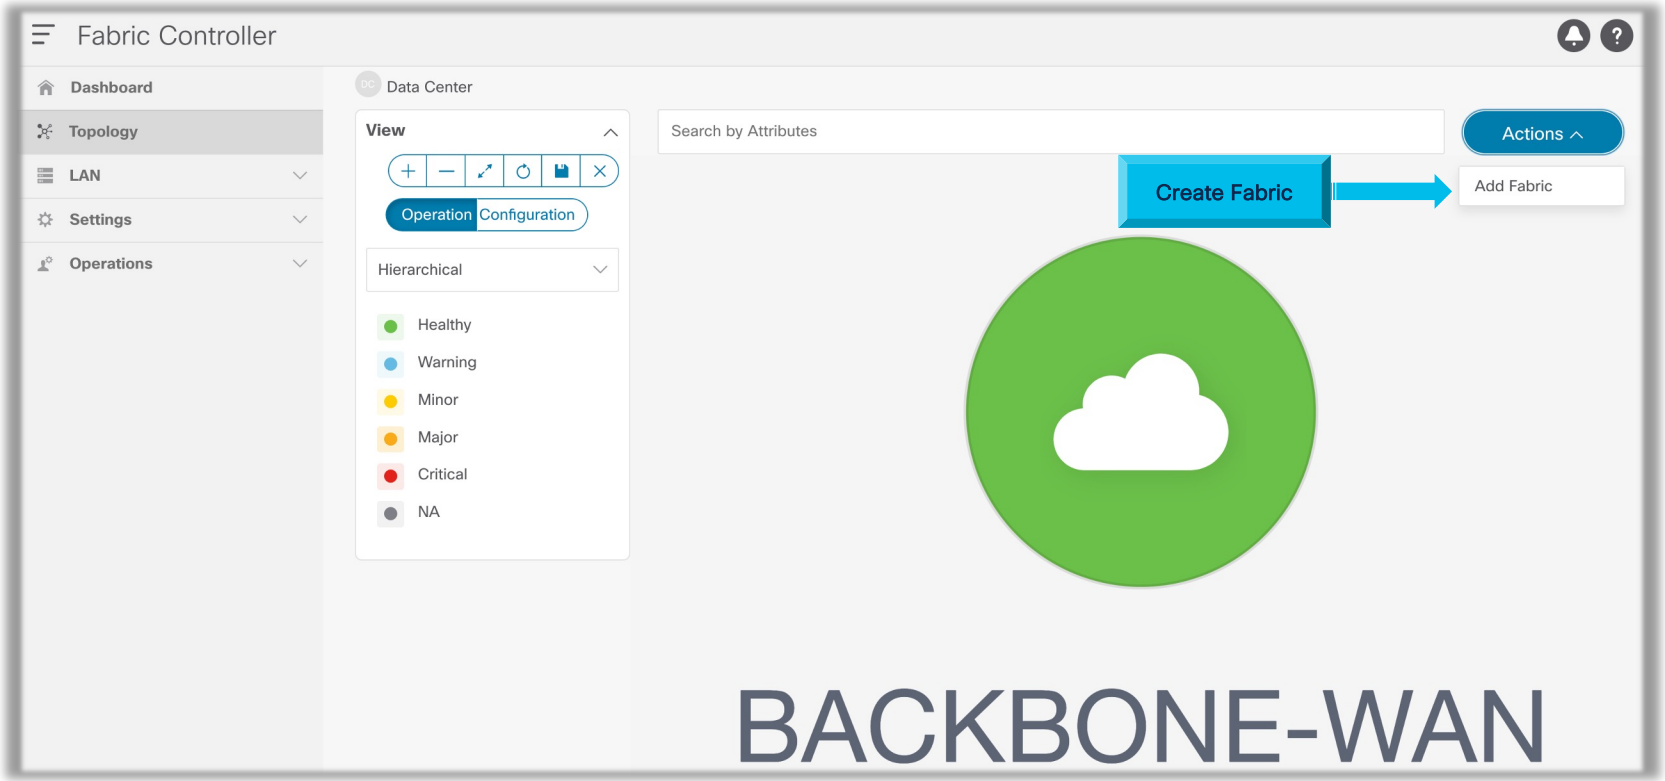

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric
(continued)

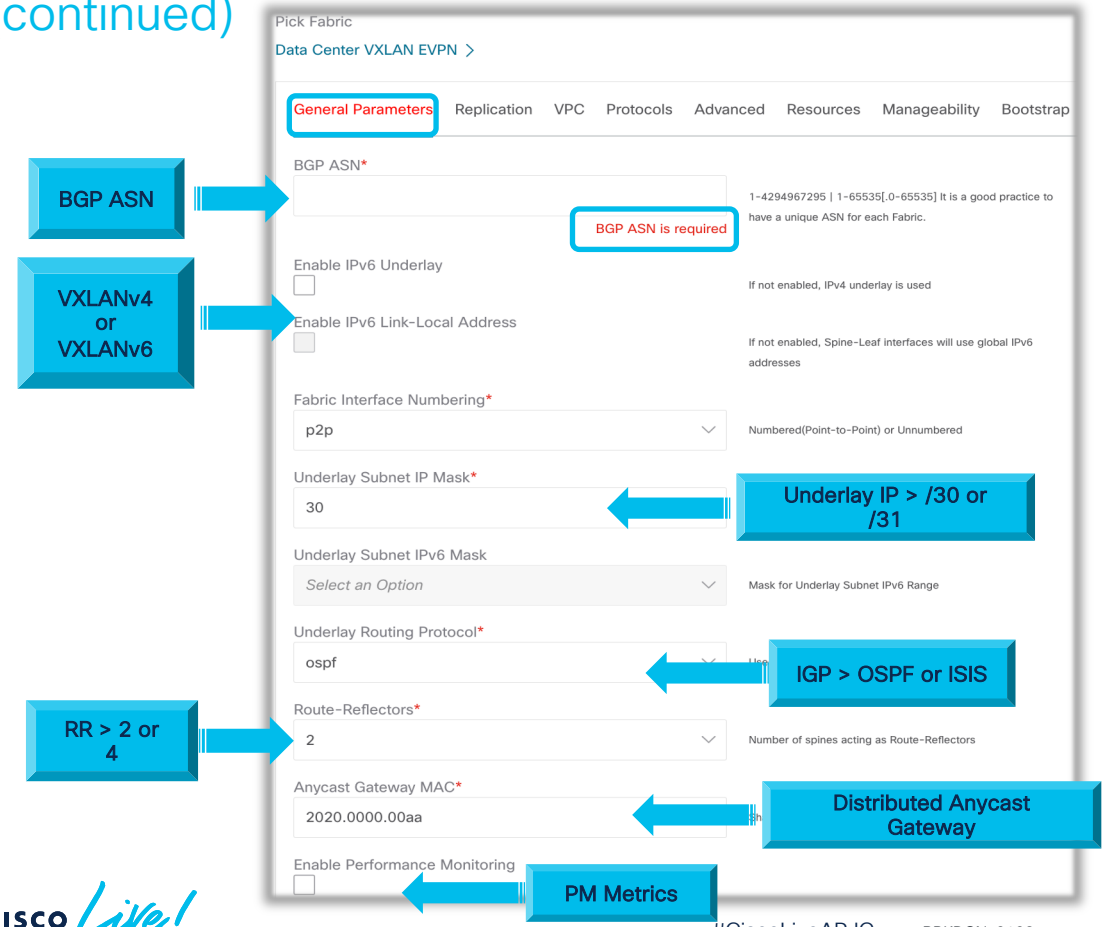

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric
(continued)

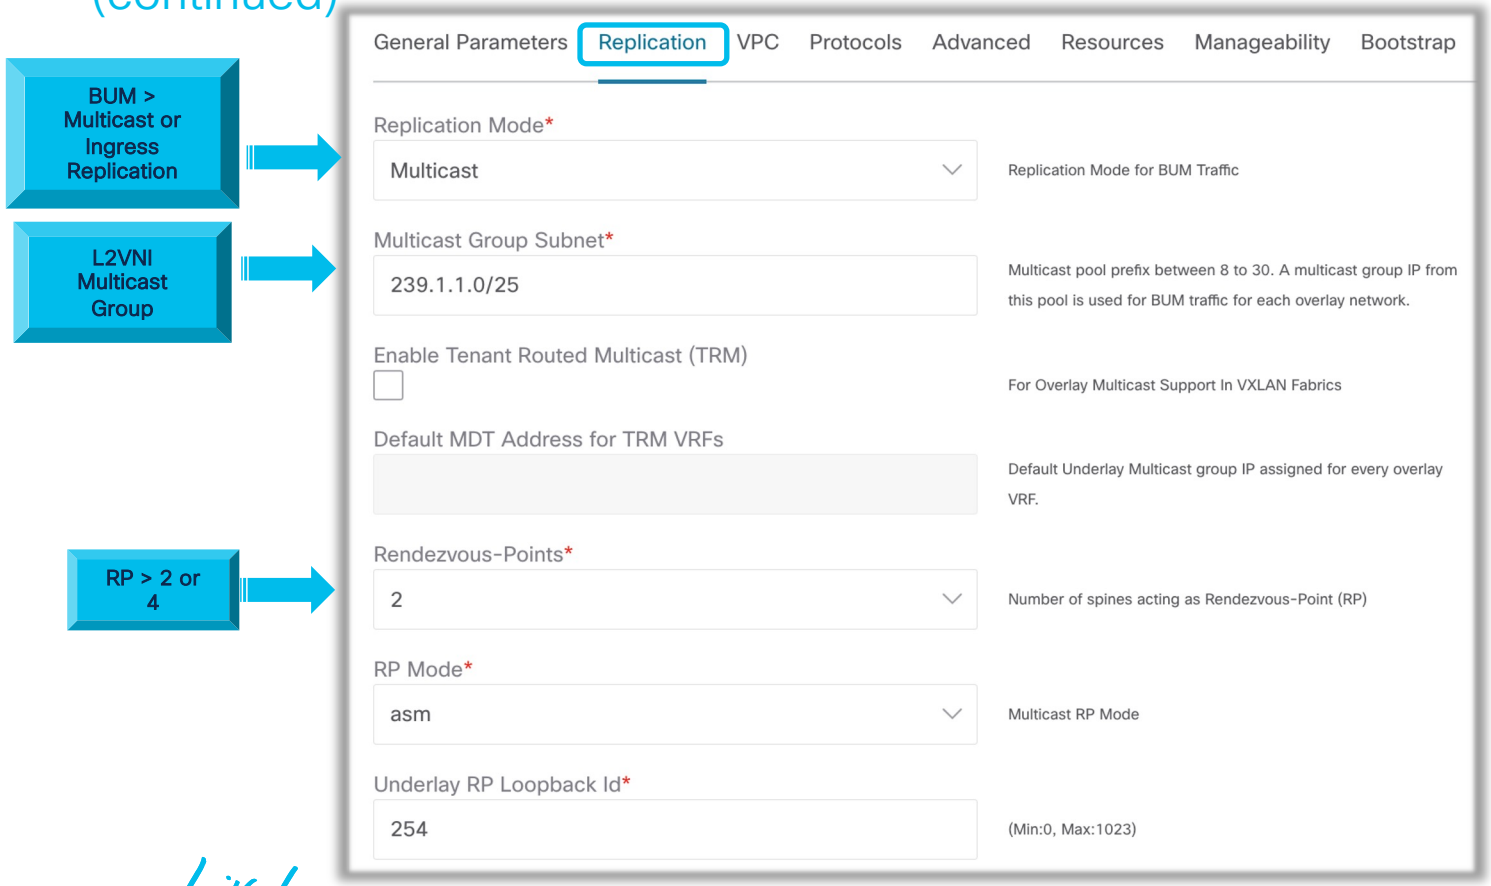

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric
(continued)

Cisco's Best Practice Configuration Templates

VXLAN Overlay Mode > CLI or Config-Profile

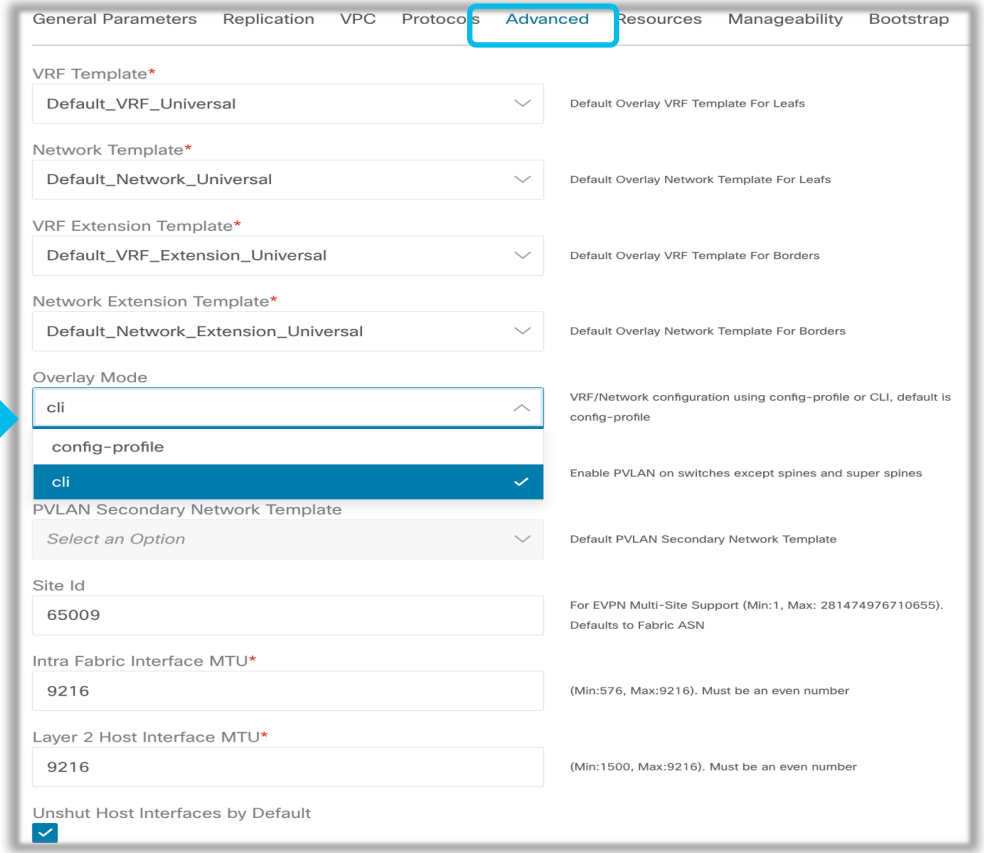

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric
(continued)

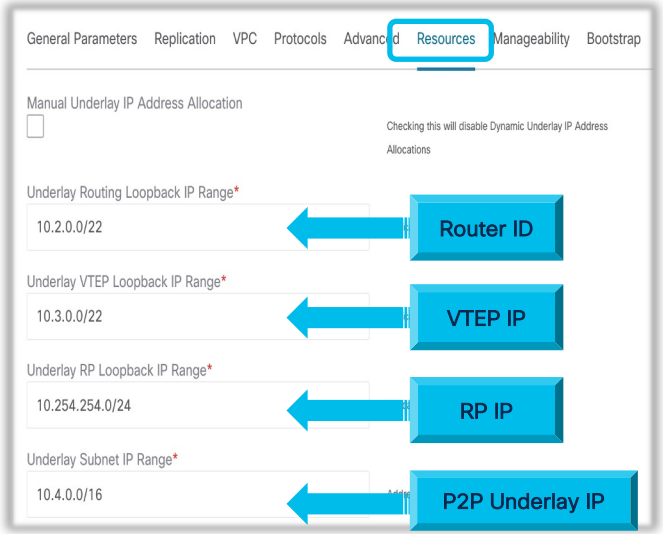

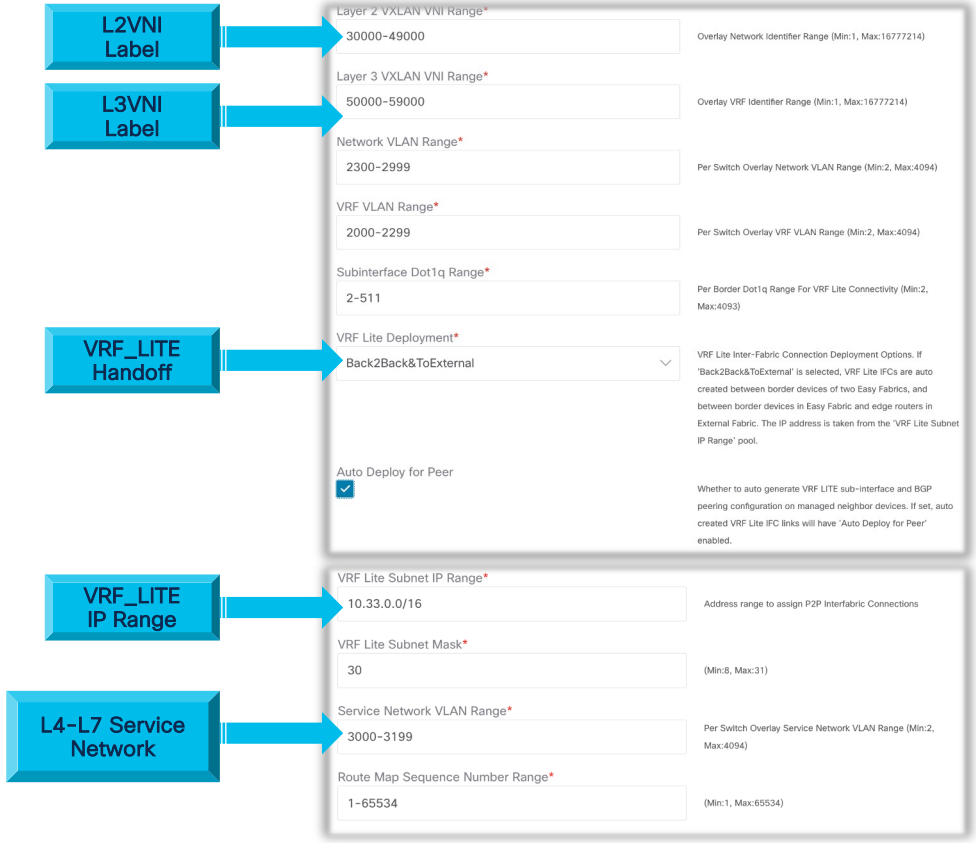

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Create a Fabric
(continued)

NDFC Built-In Bootstrap POAP Services. Supports OOB and In band POAP

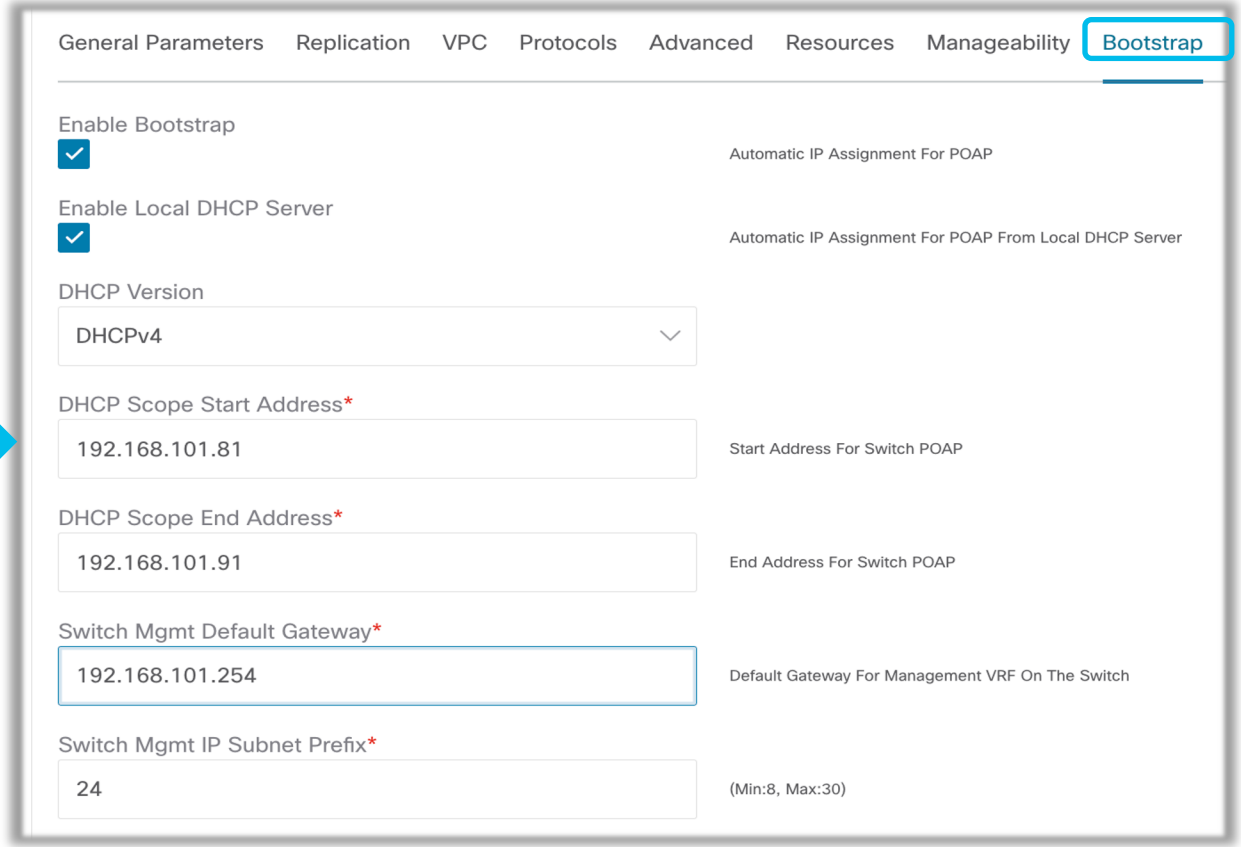

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 2 -> Add Switches

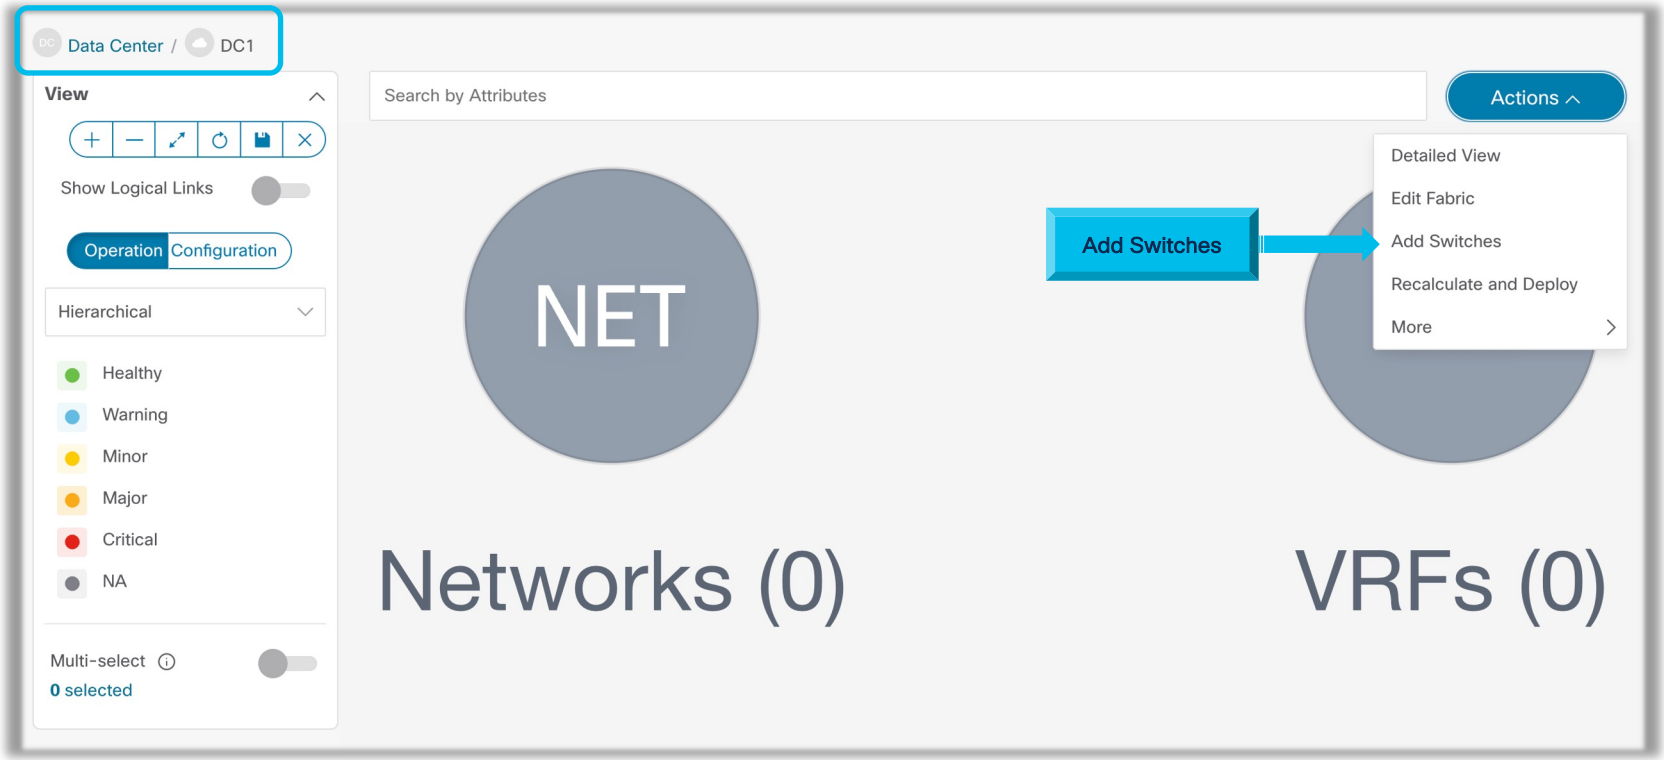

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 2 -> Add Switches
(continued)

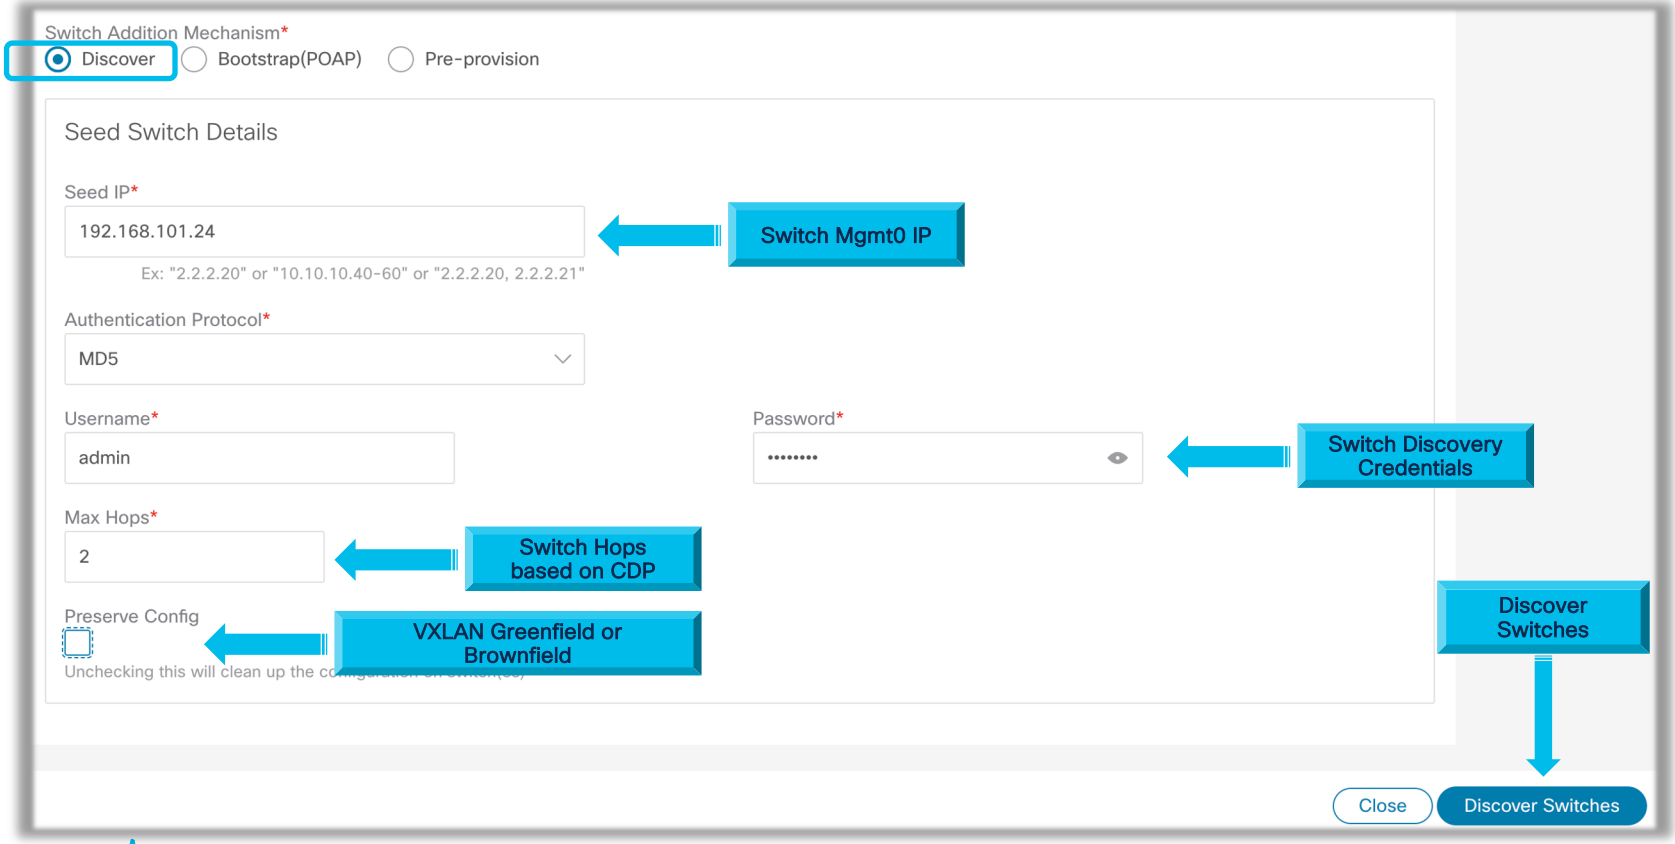

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 2 -> Add Switches
(continued)

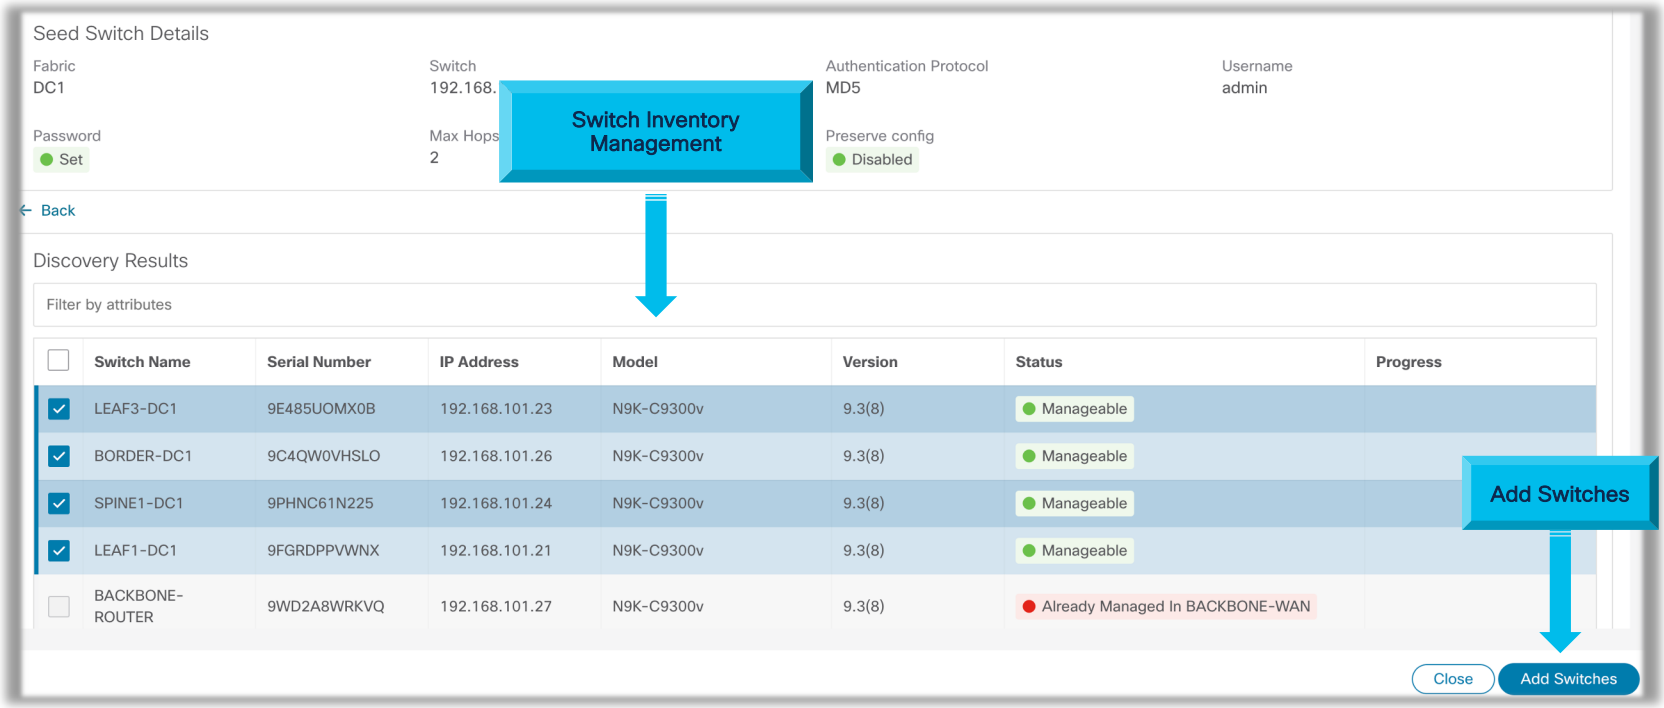

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Set Role

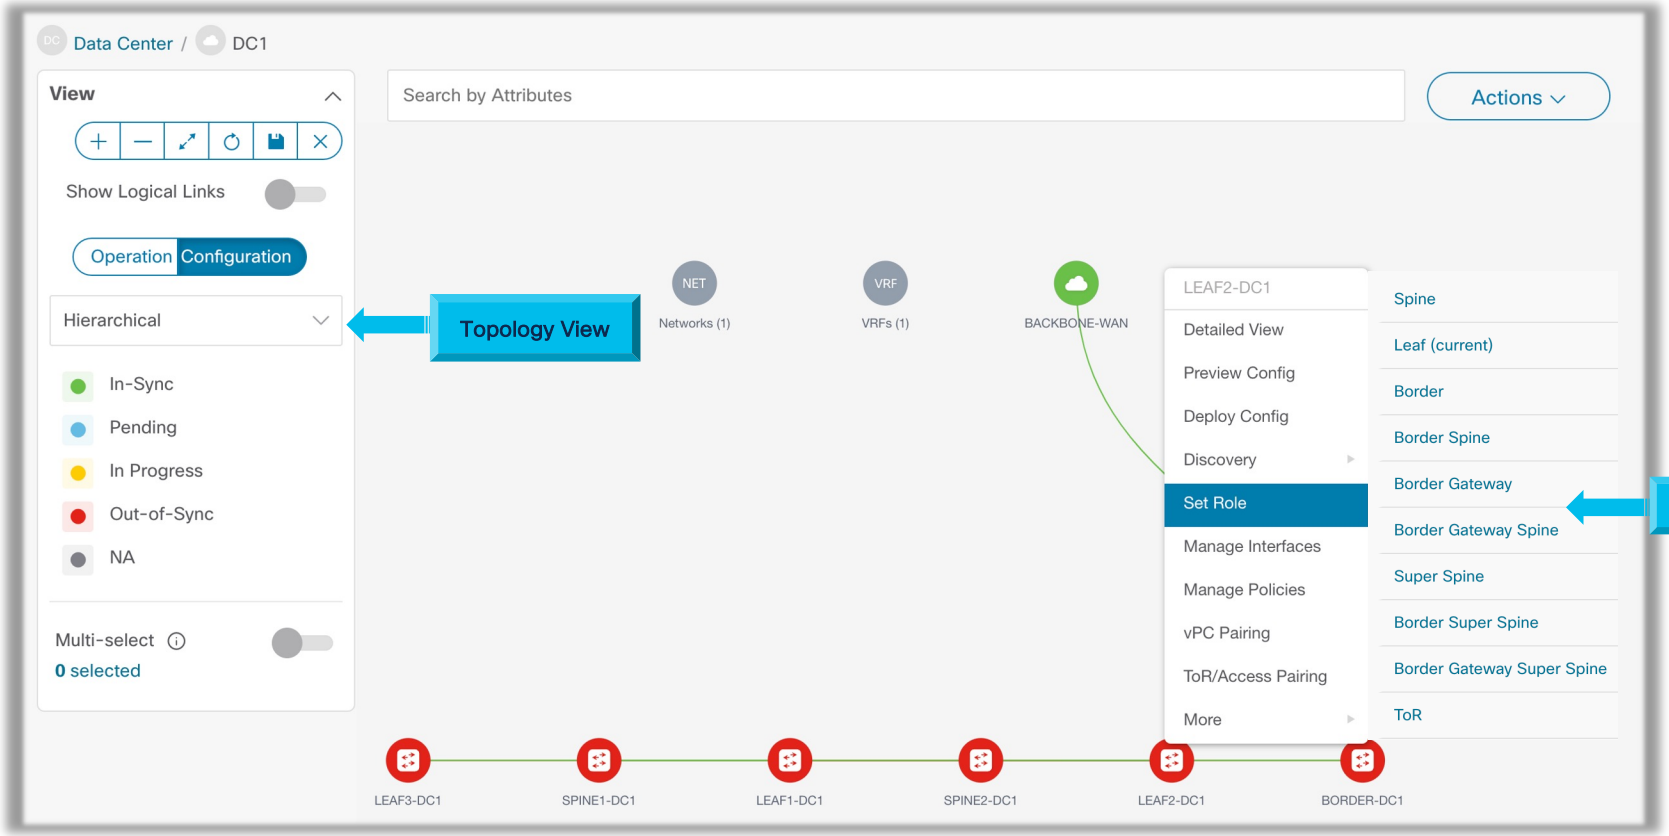

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 4 -> VPC Pairing
(optional)

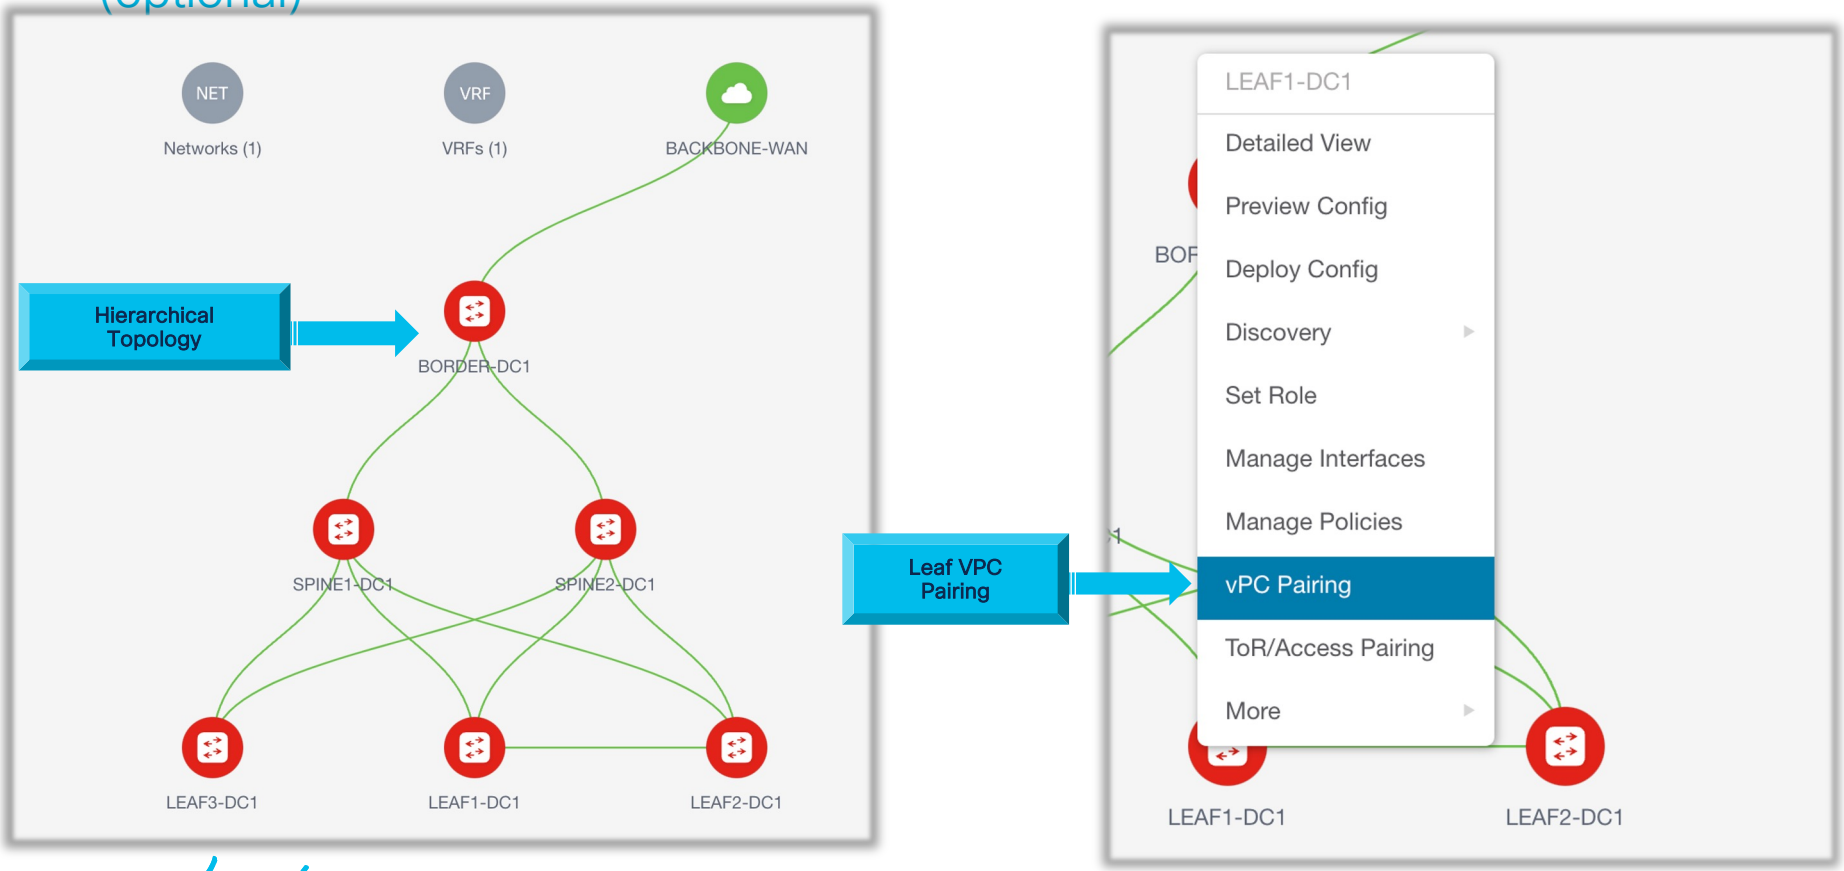

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 4 -> VPC Pairing
(optional)

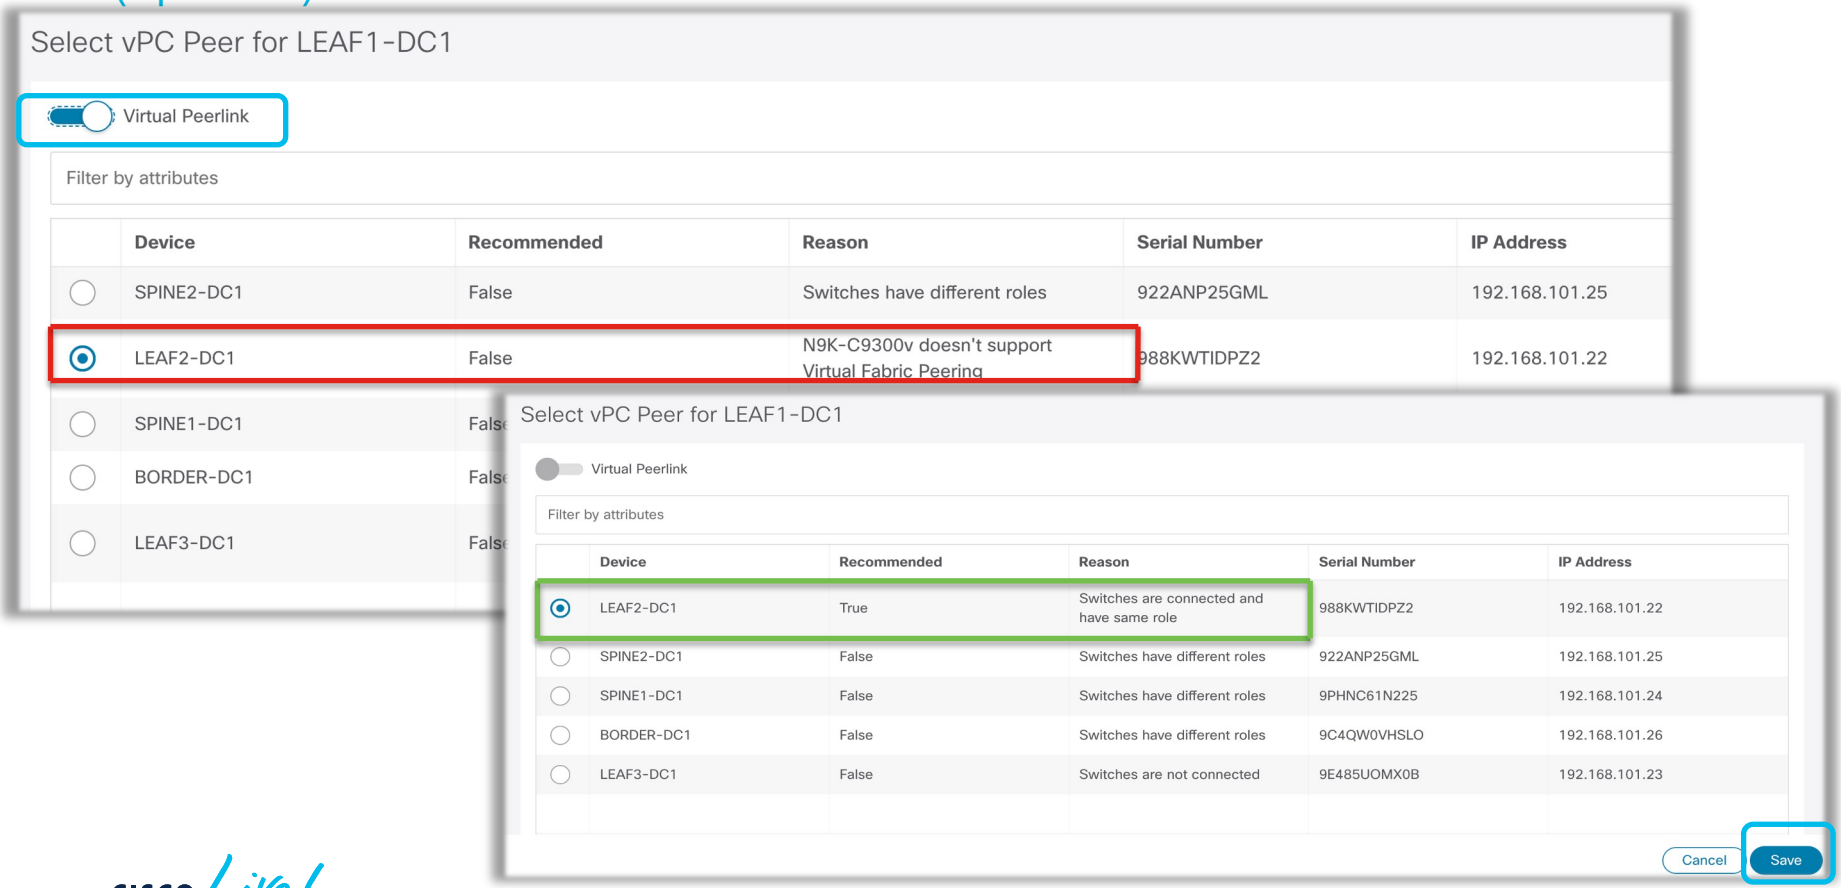

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 5 -> Recalculate and Deploy

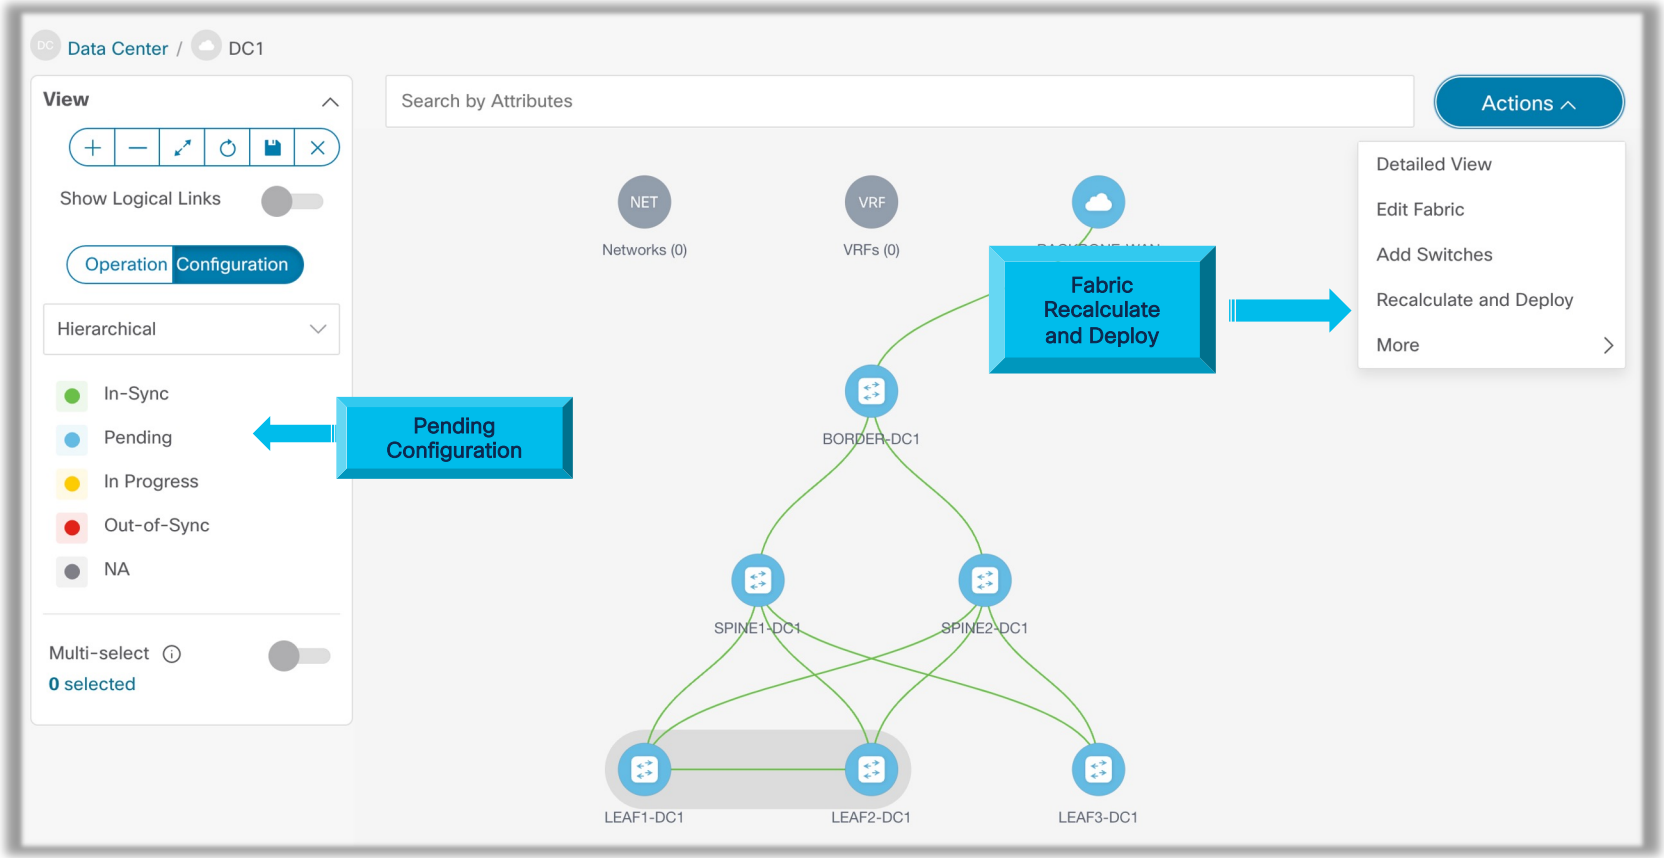

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 5 -> Recalculate and Deploy
(continued)

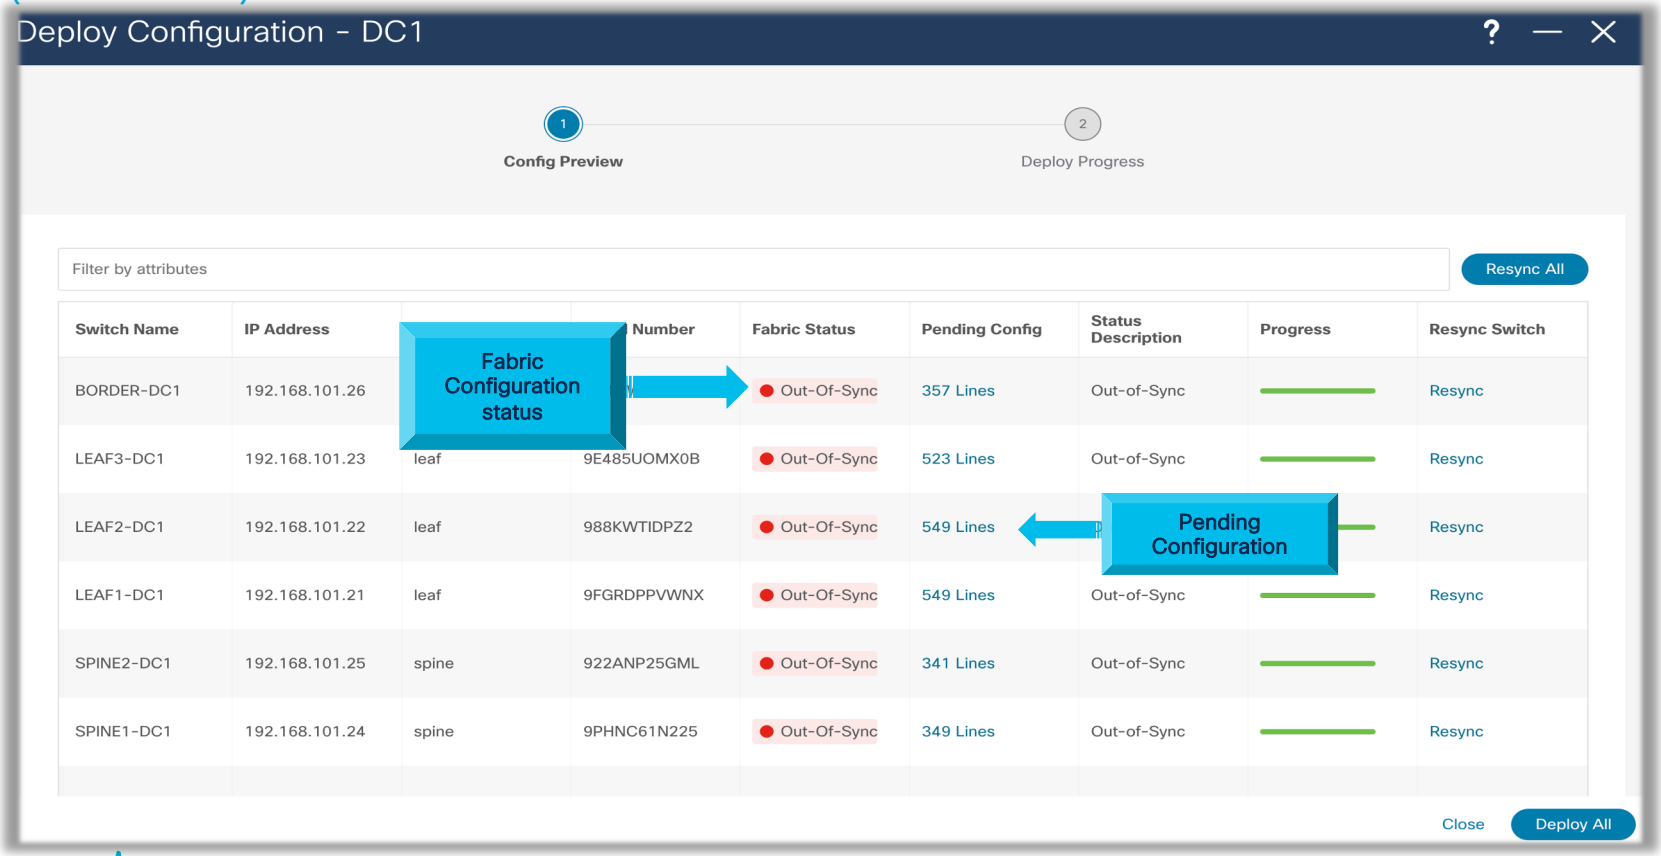

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 5 -> Recalculate and Deploy
(continued)

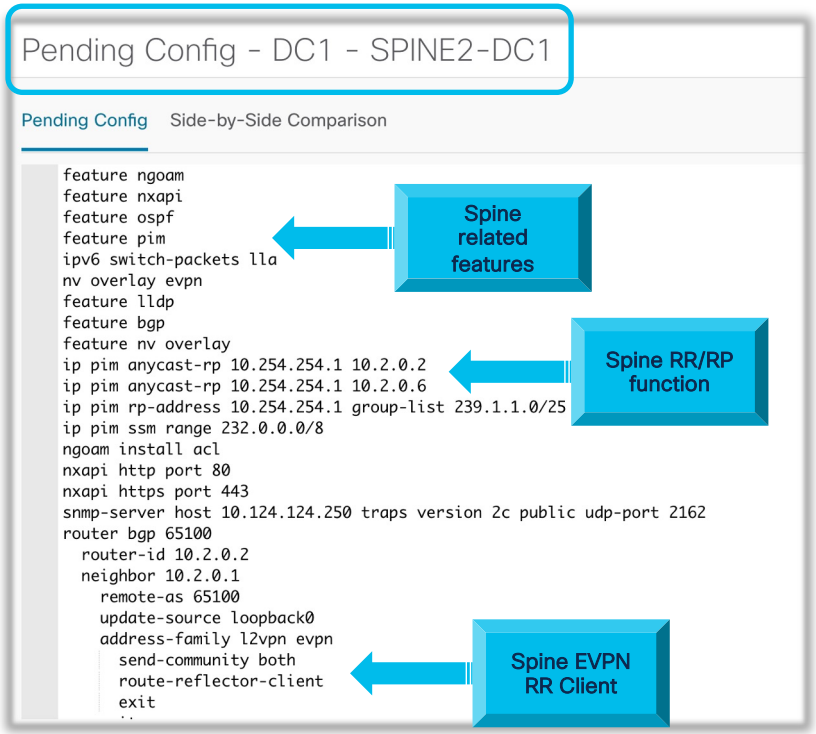

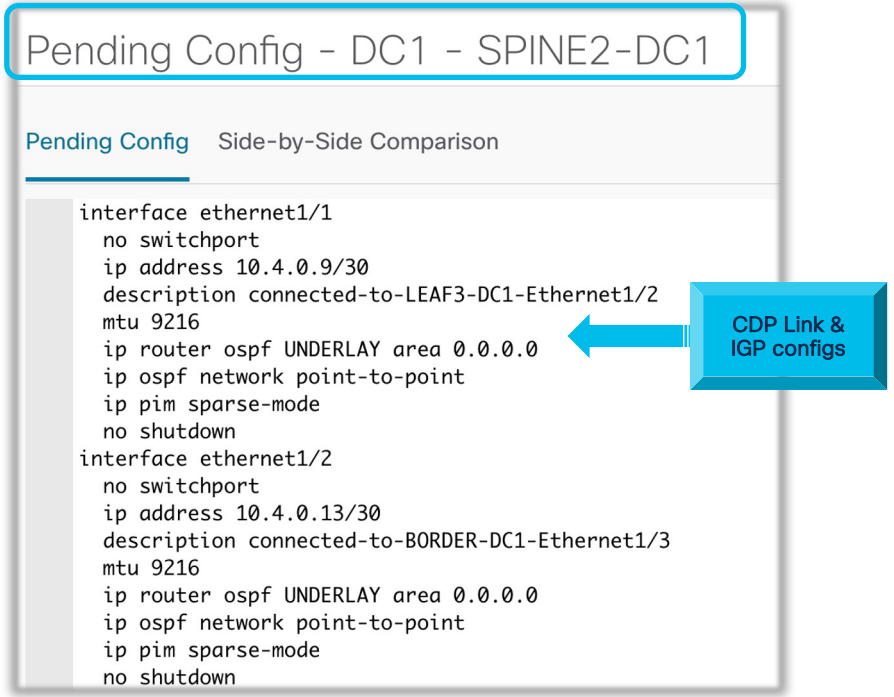

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 5 -> Recalculate and Deploy
(continued)

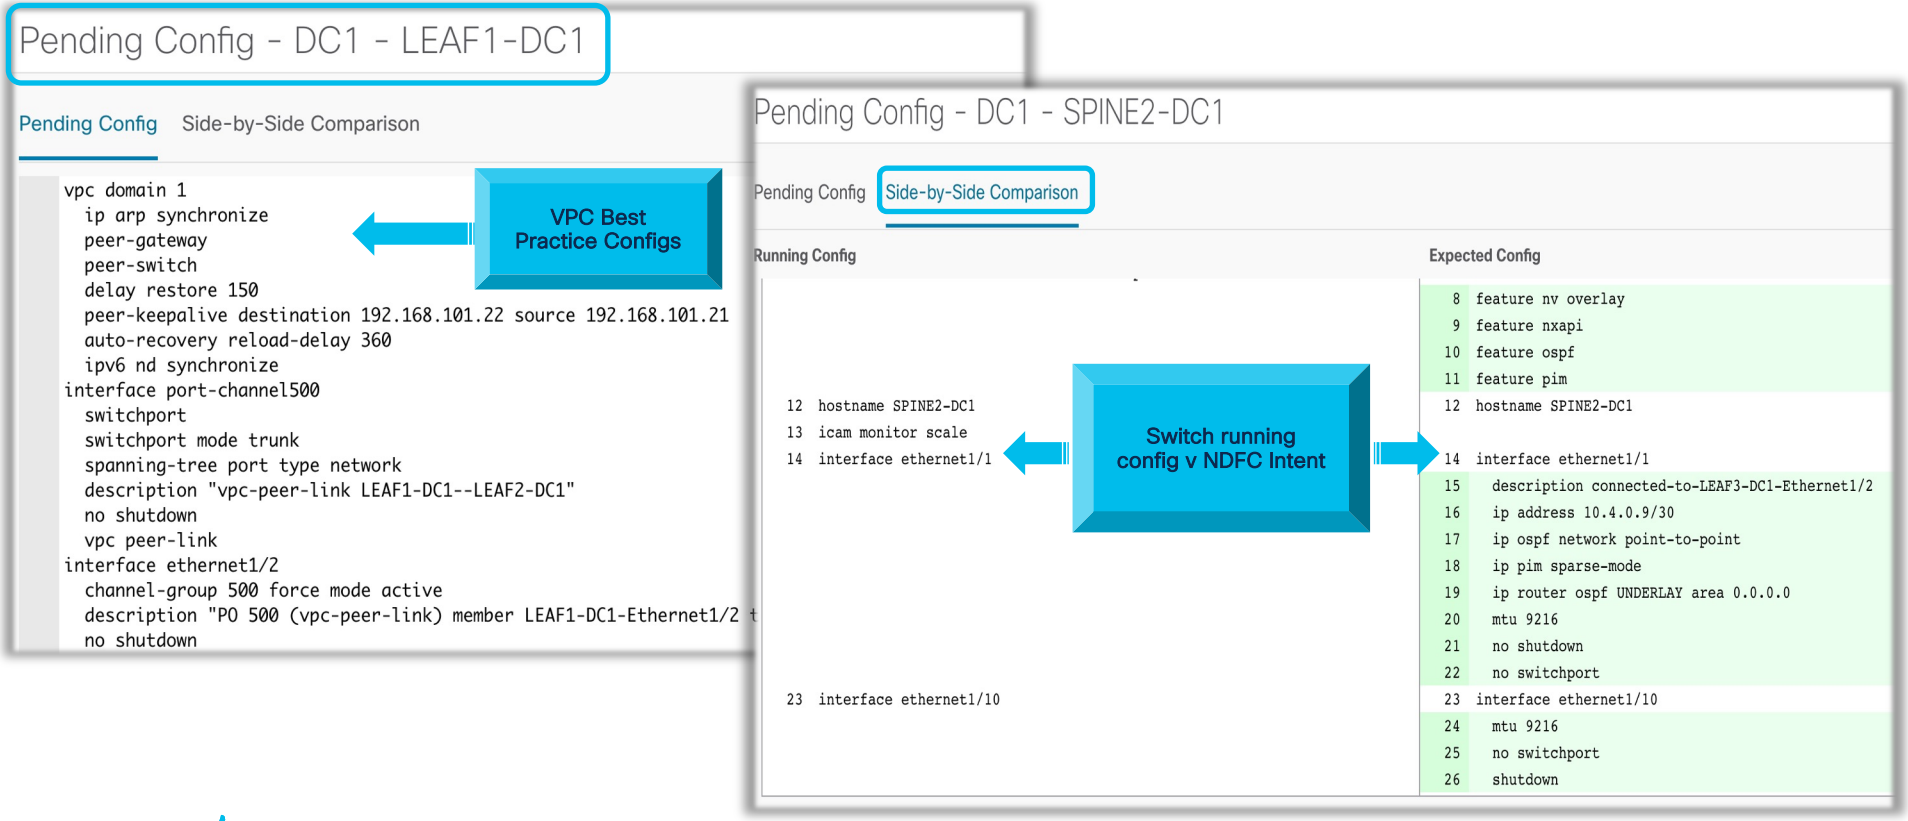

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

NDFC VXLAN EVPN Topology View

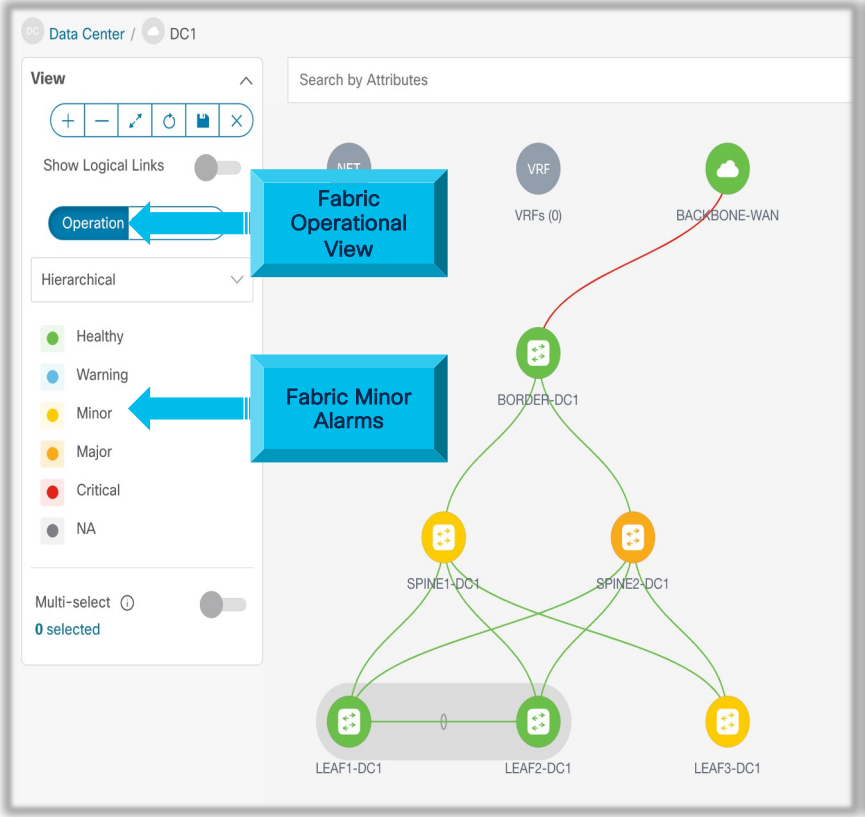

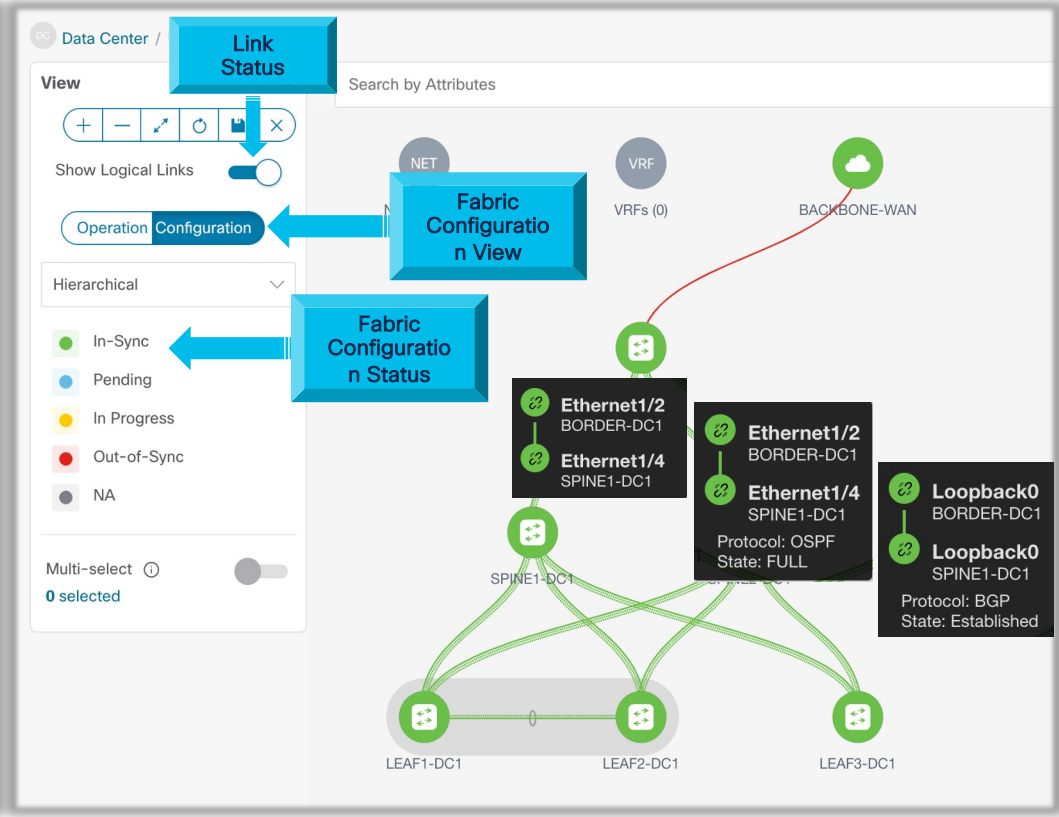

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 & 2 -> Create Individual Fabrics

Set appropriate roles

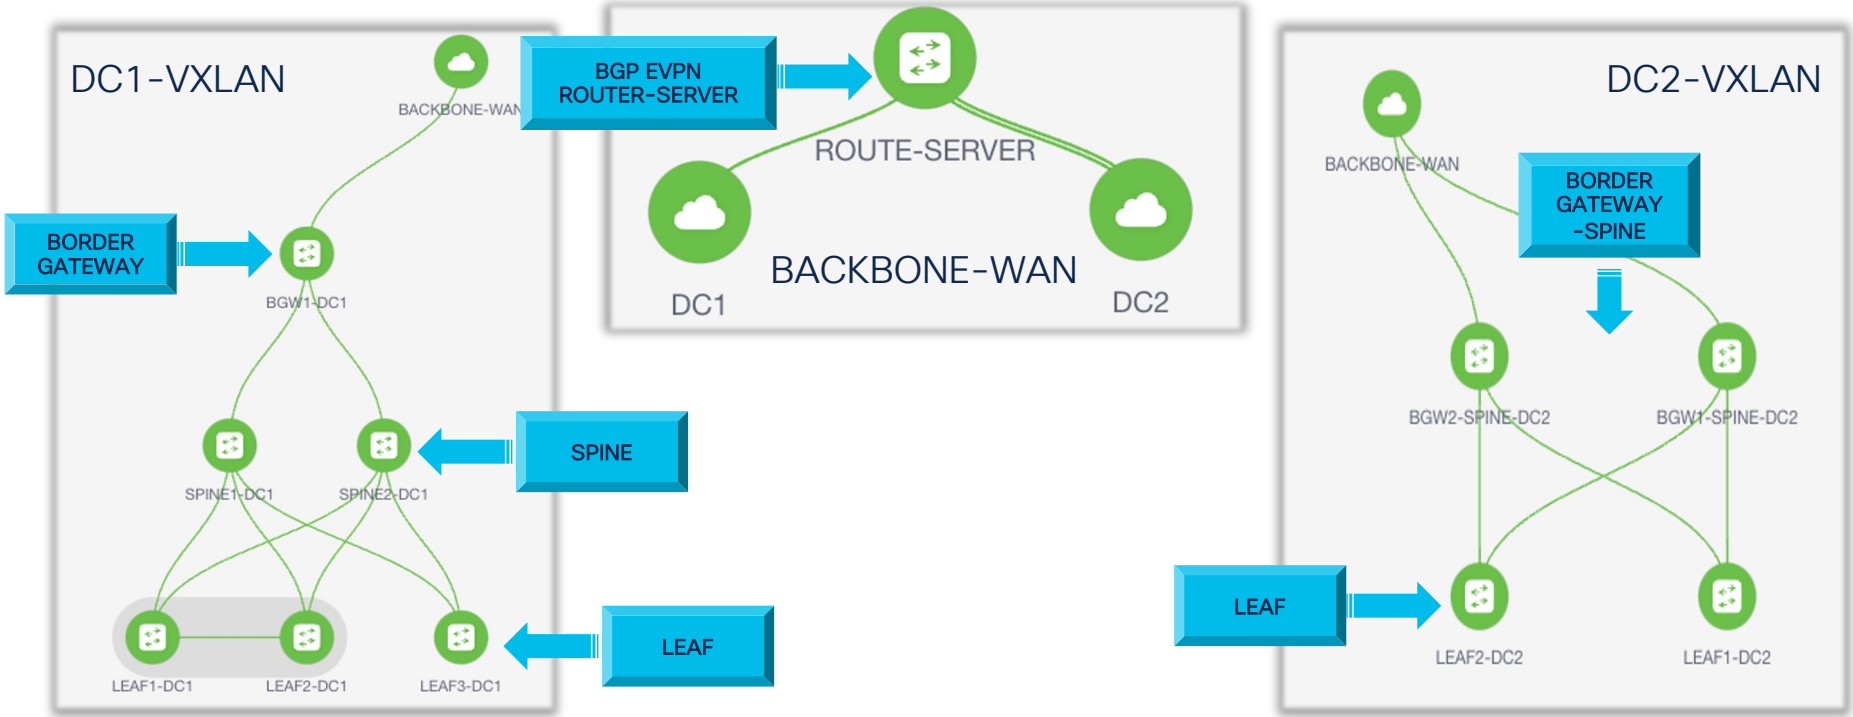

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Create VXLAN EVPN Multi-Site Fabric

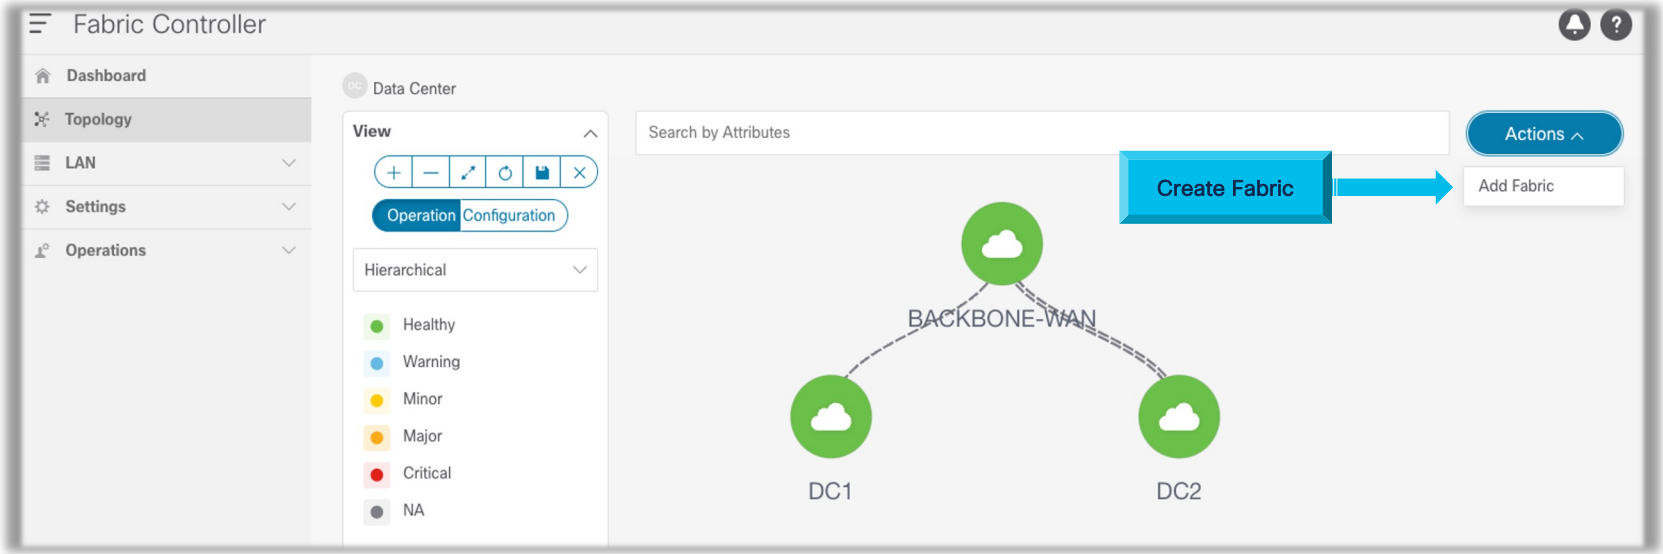

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Create VXLAN EVPN Multi-Site Fabric (continued)

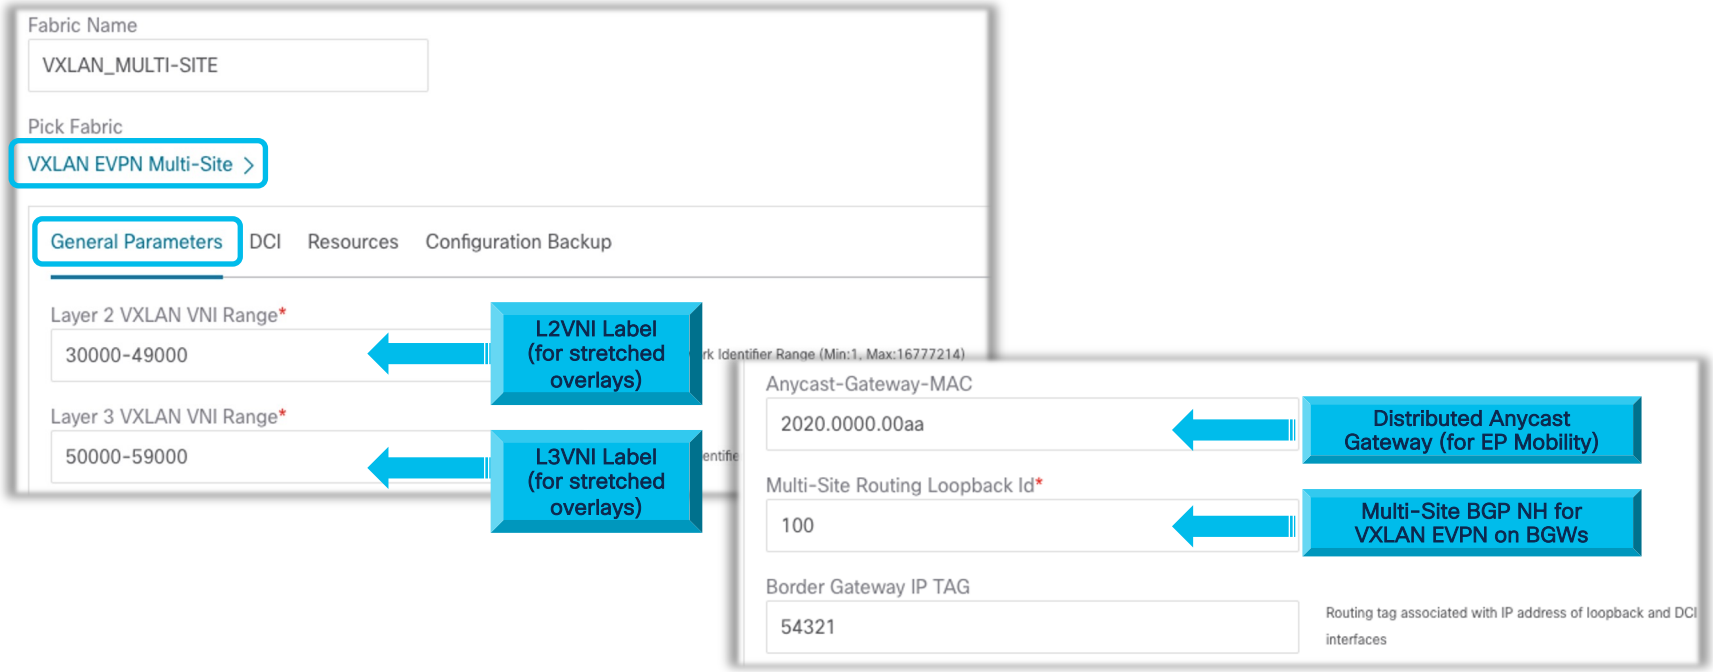

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Create VXLAN EVPN Multi-Site Fabric (continued)

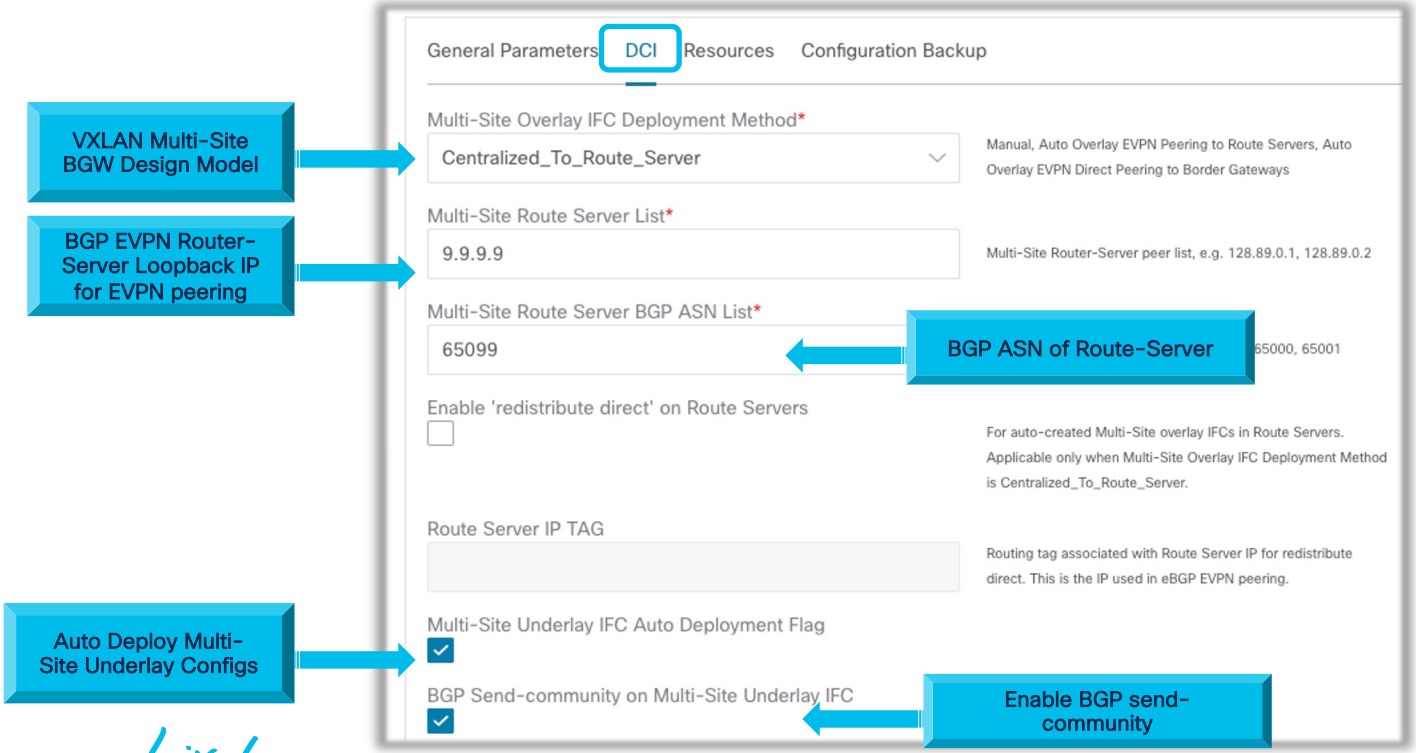

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Create VXLAN EVPN Multi-Site Fabric (continued)

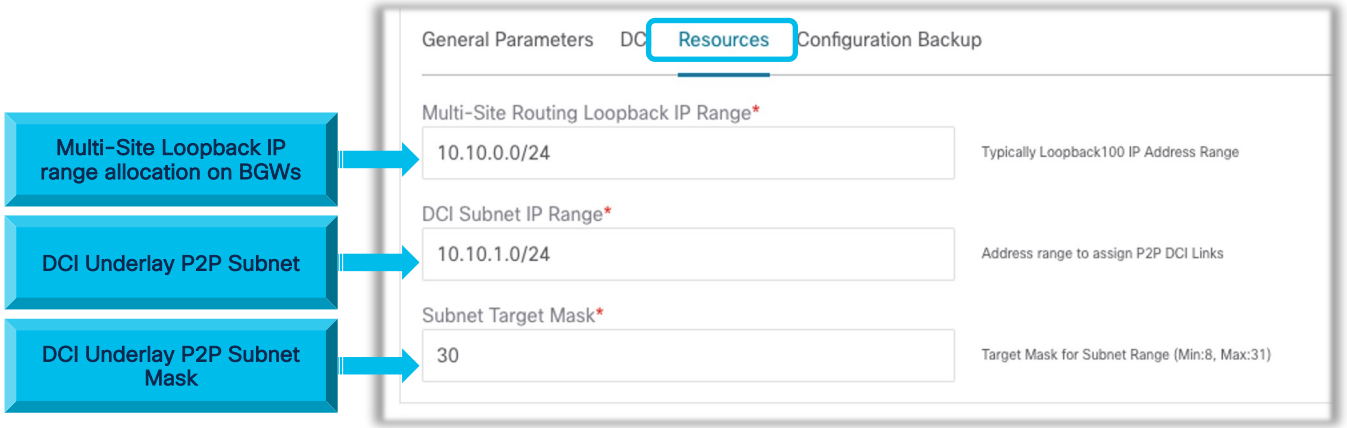

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Create VXLAN EVPN Multi-Site Fabric
(continued)

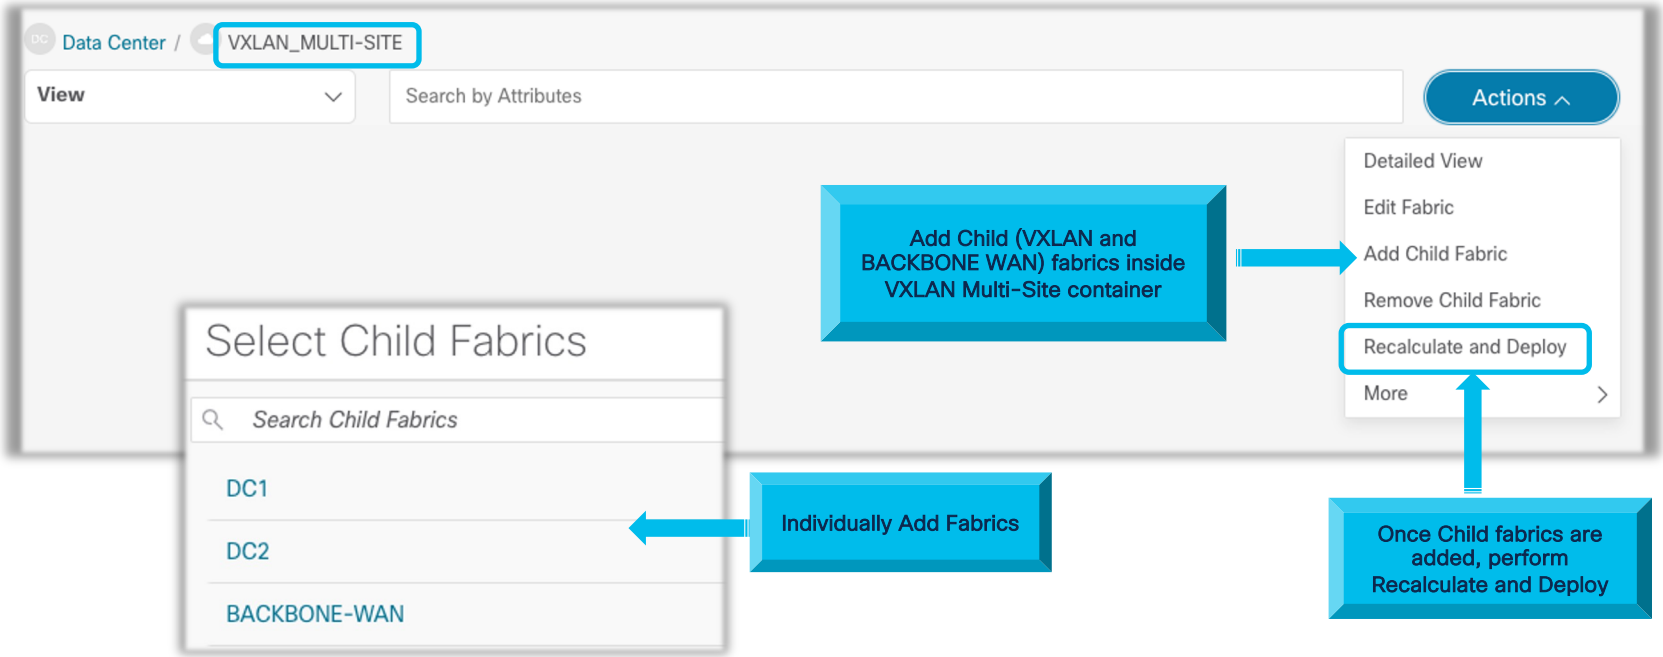

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 4 -> Recalculate & Deploy

In VXLAN EVPN Multi-Site Fabric

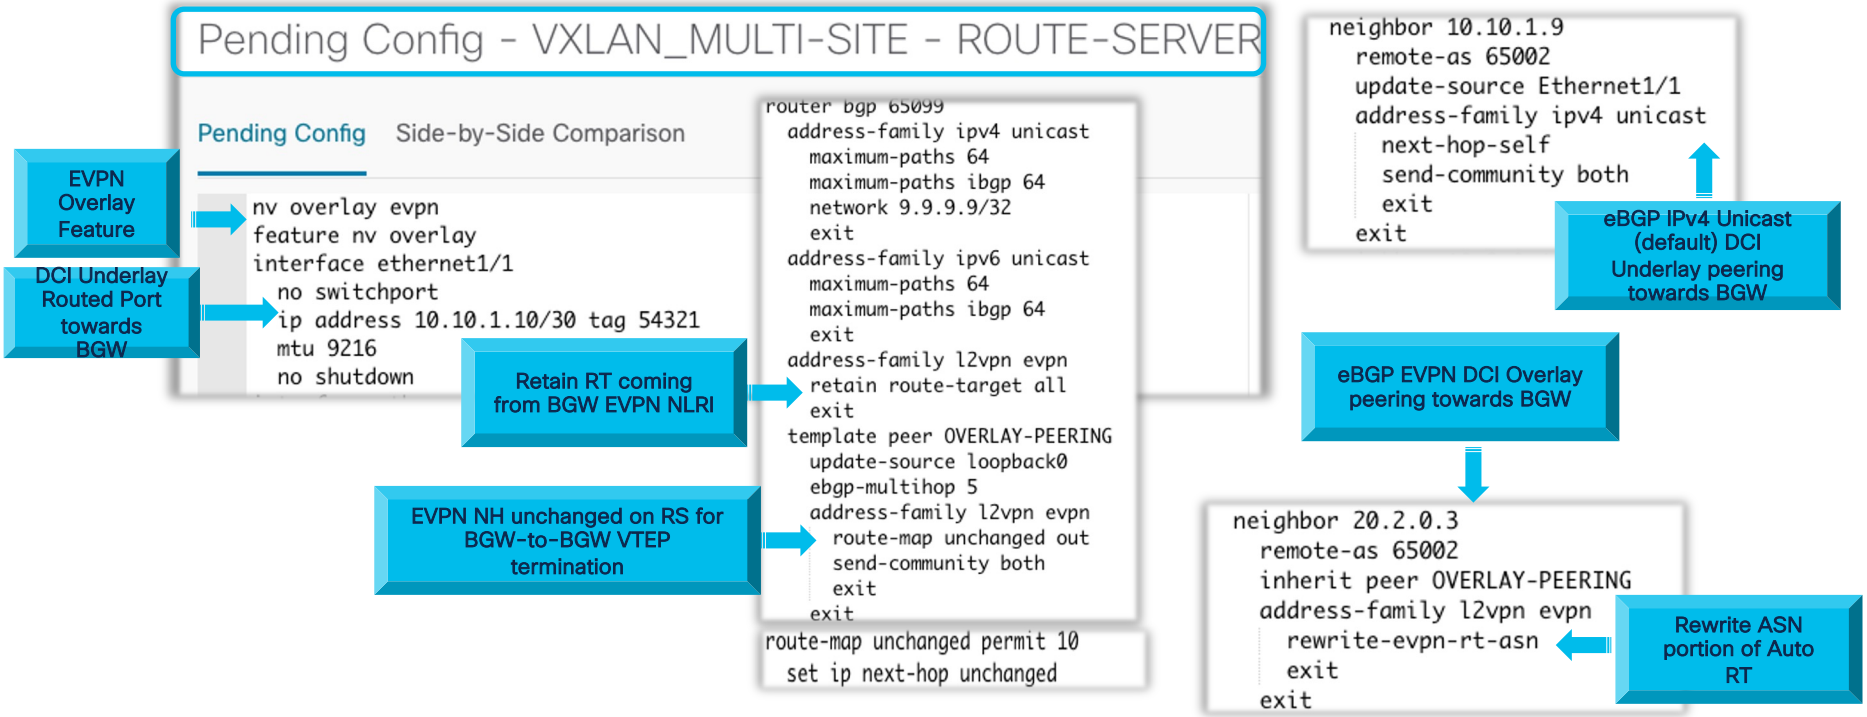

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 4 -> Recalculate & Deploy

In VXLAN EVPN Multi-Site Fabric

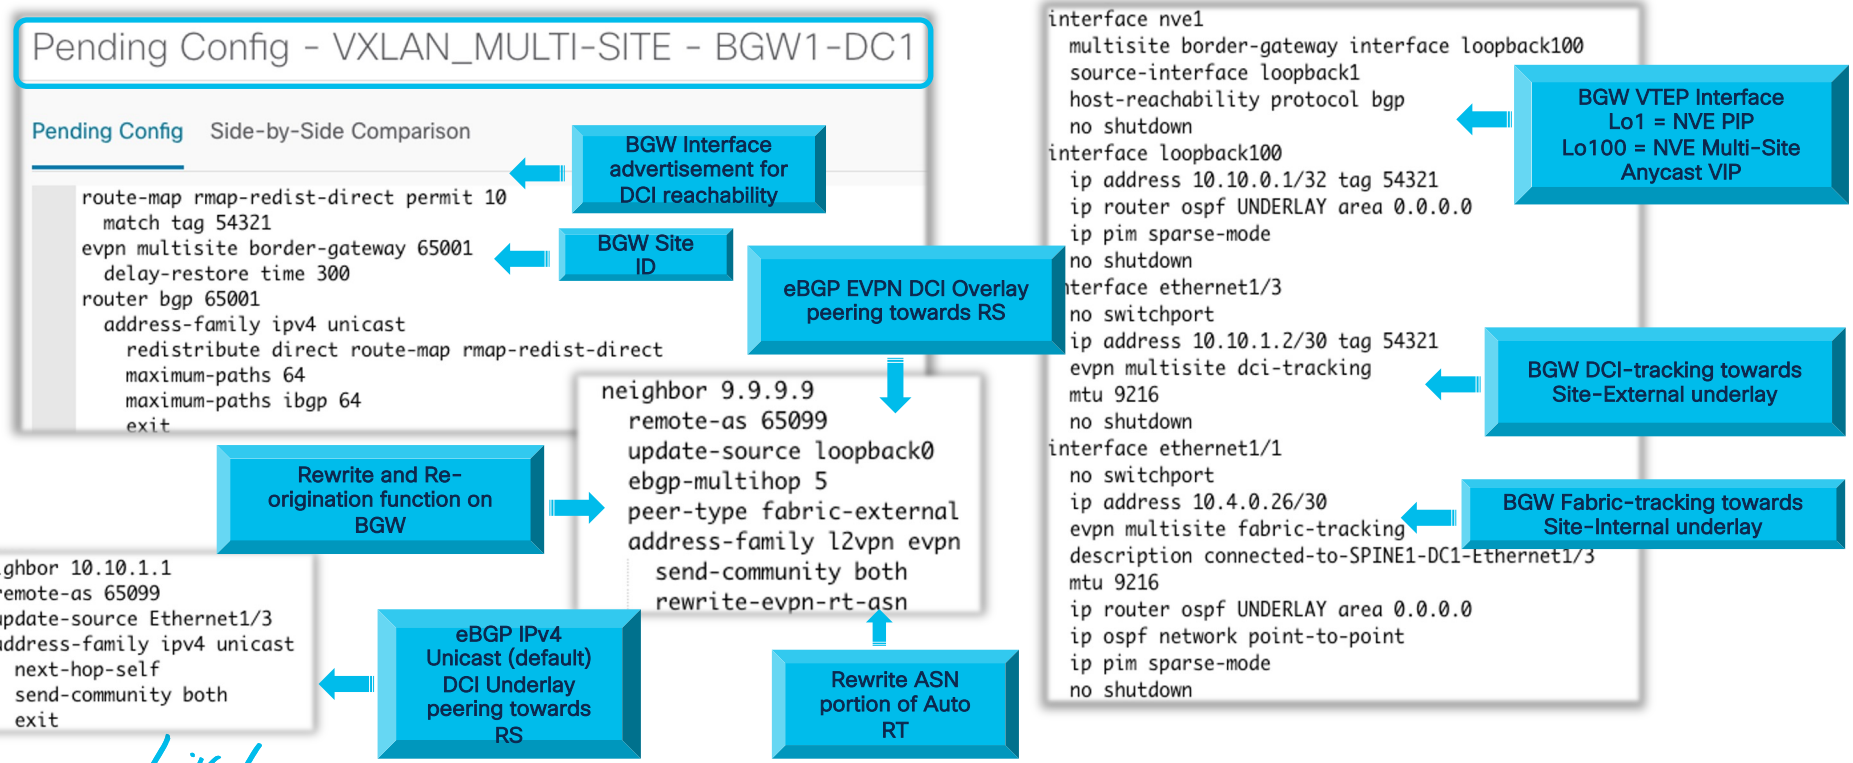

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

NDFC VXLAN EVPN Multi-Site Topology View

Site-External DCI

Site-Internal Fabric

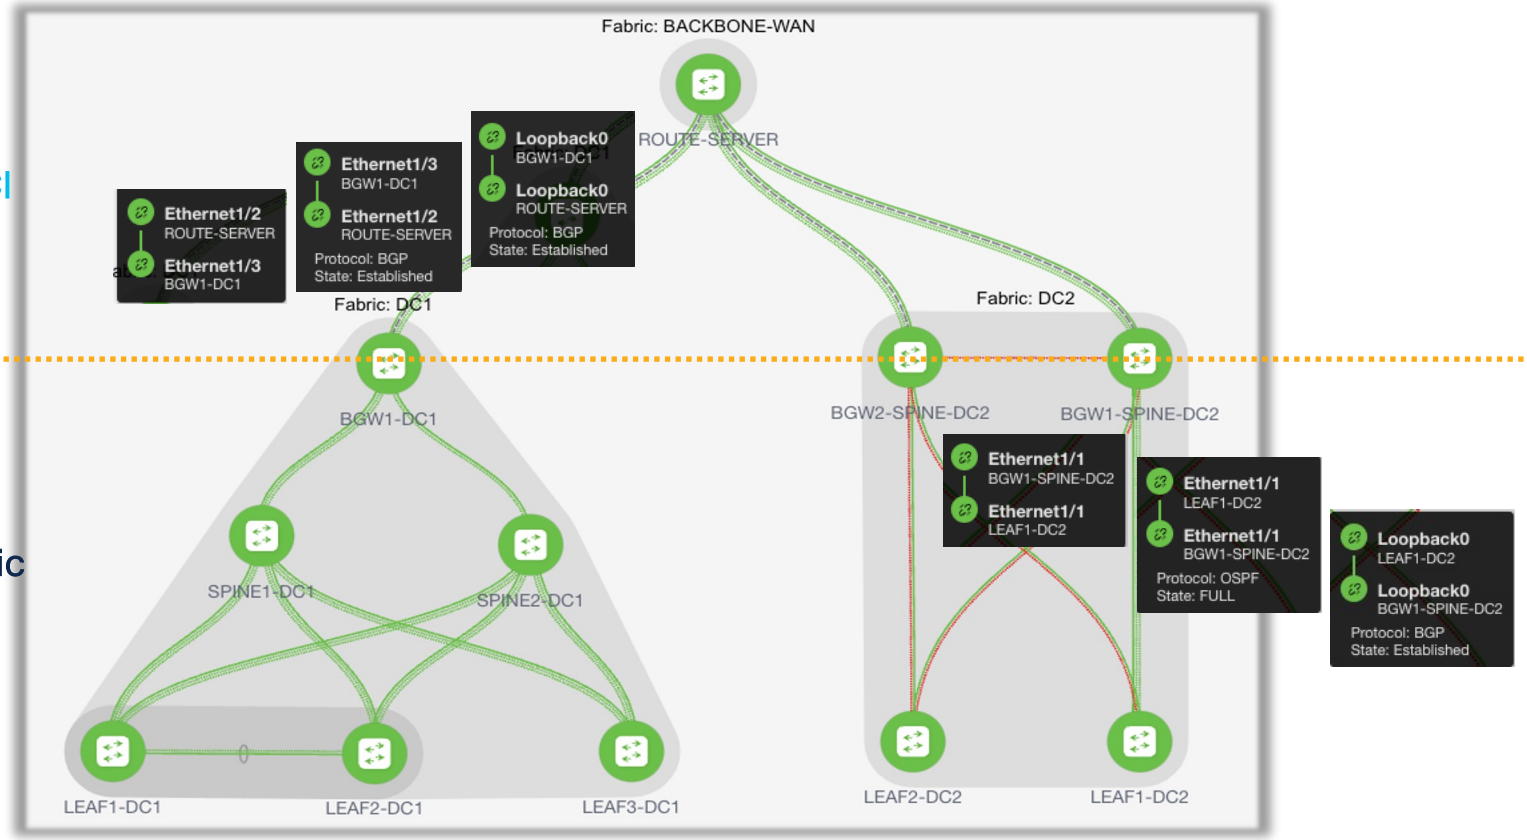

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Deploy VXLAN EVPN Overlays

in Multi-Site Architecture

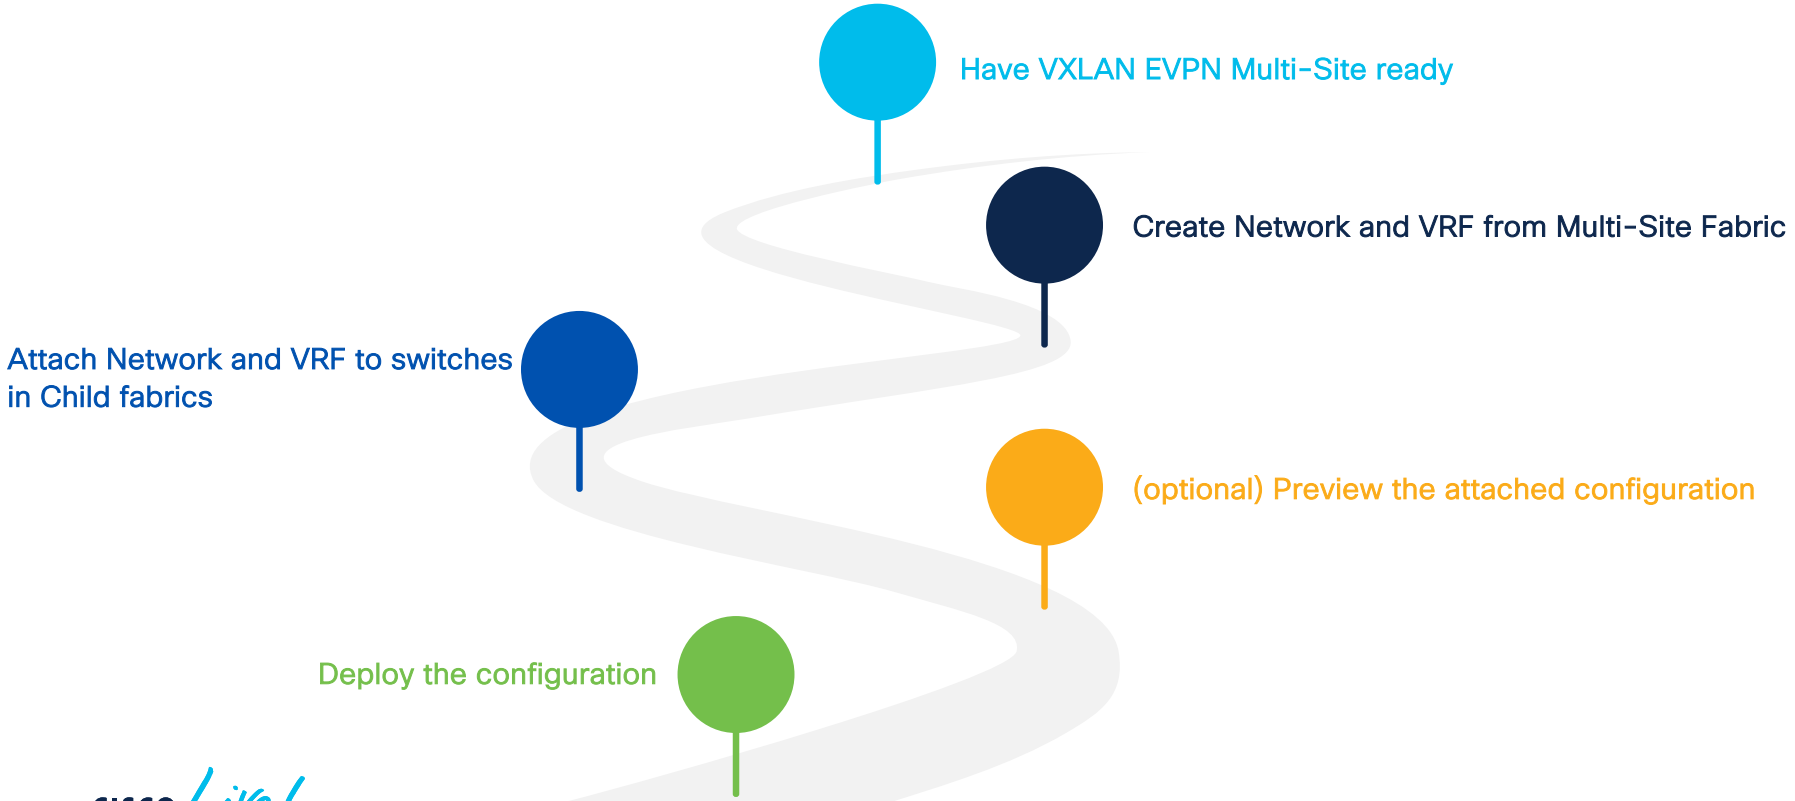

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 1 -> Navigate to VXLAN Multi-Site Fabric

Select Multi-Site Fabric

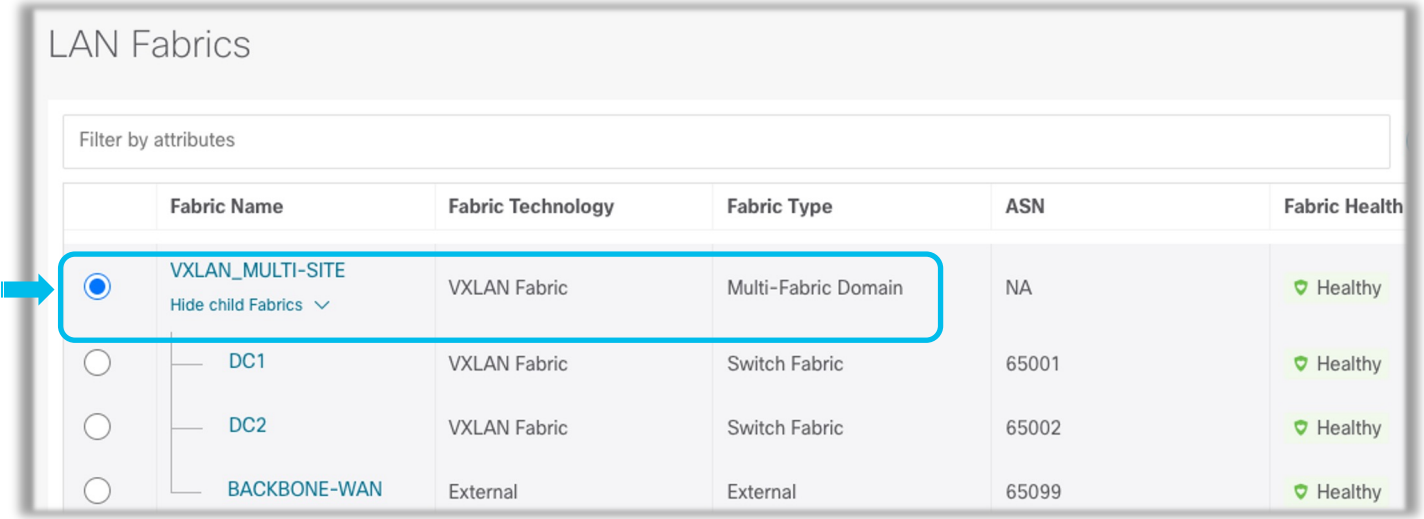

*UI Navigation: LAN Fabrics > Select Fabric > Fabric Overview > Networks > Create

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 2 -> Create VRF and Network

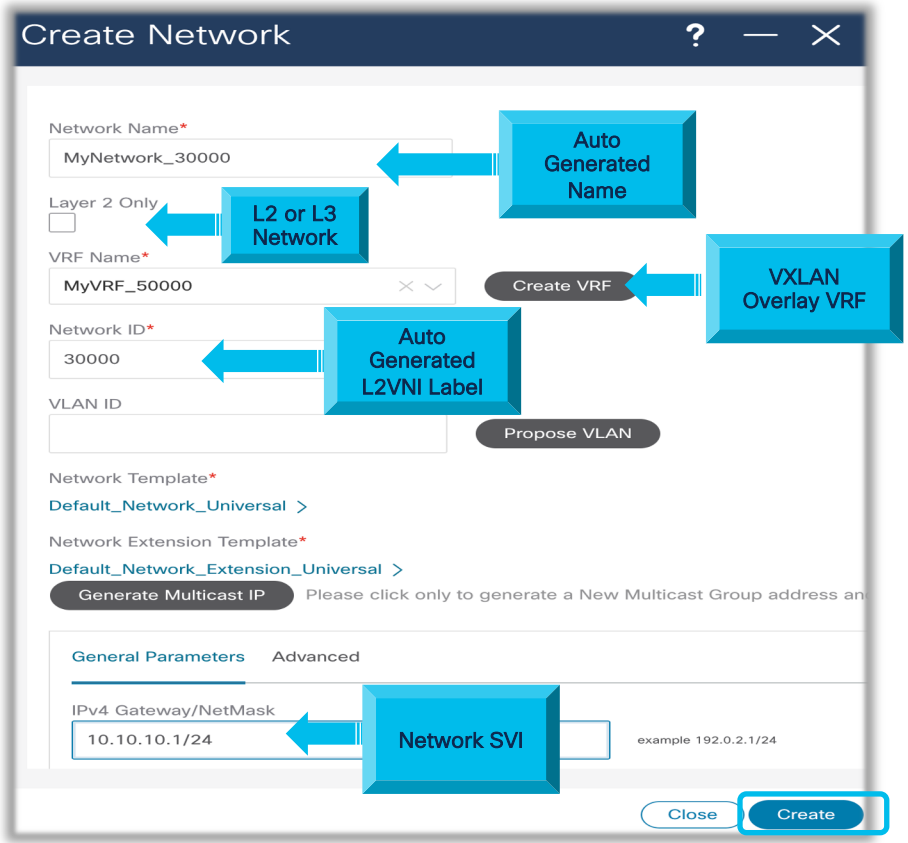

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Attach VRF and Network

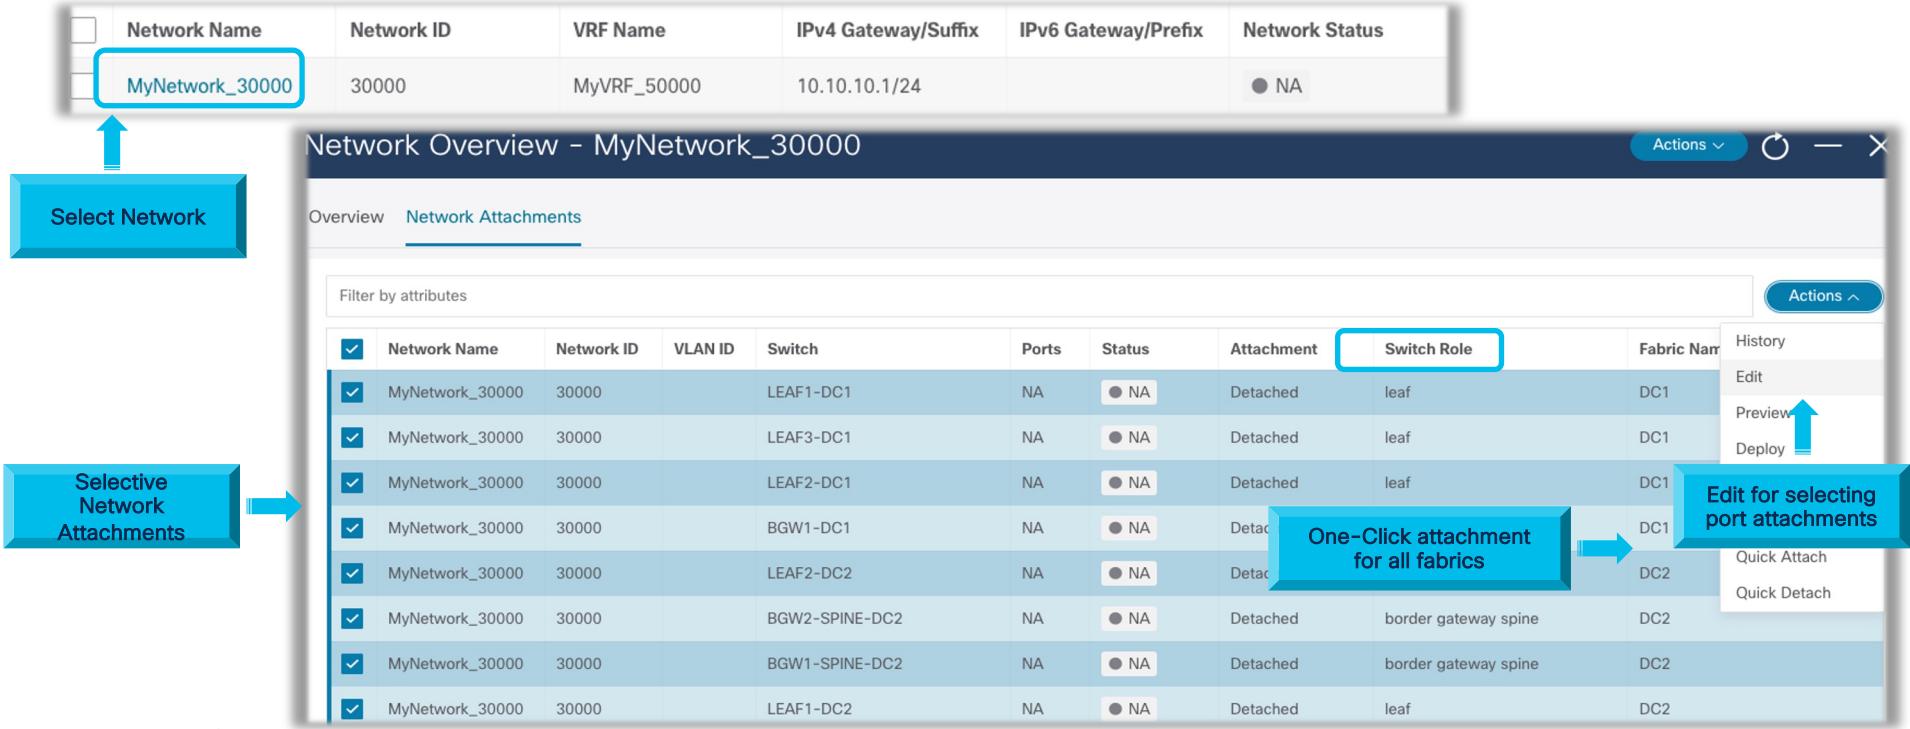

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 3 -> Attach VRF and Network

(continued)

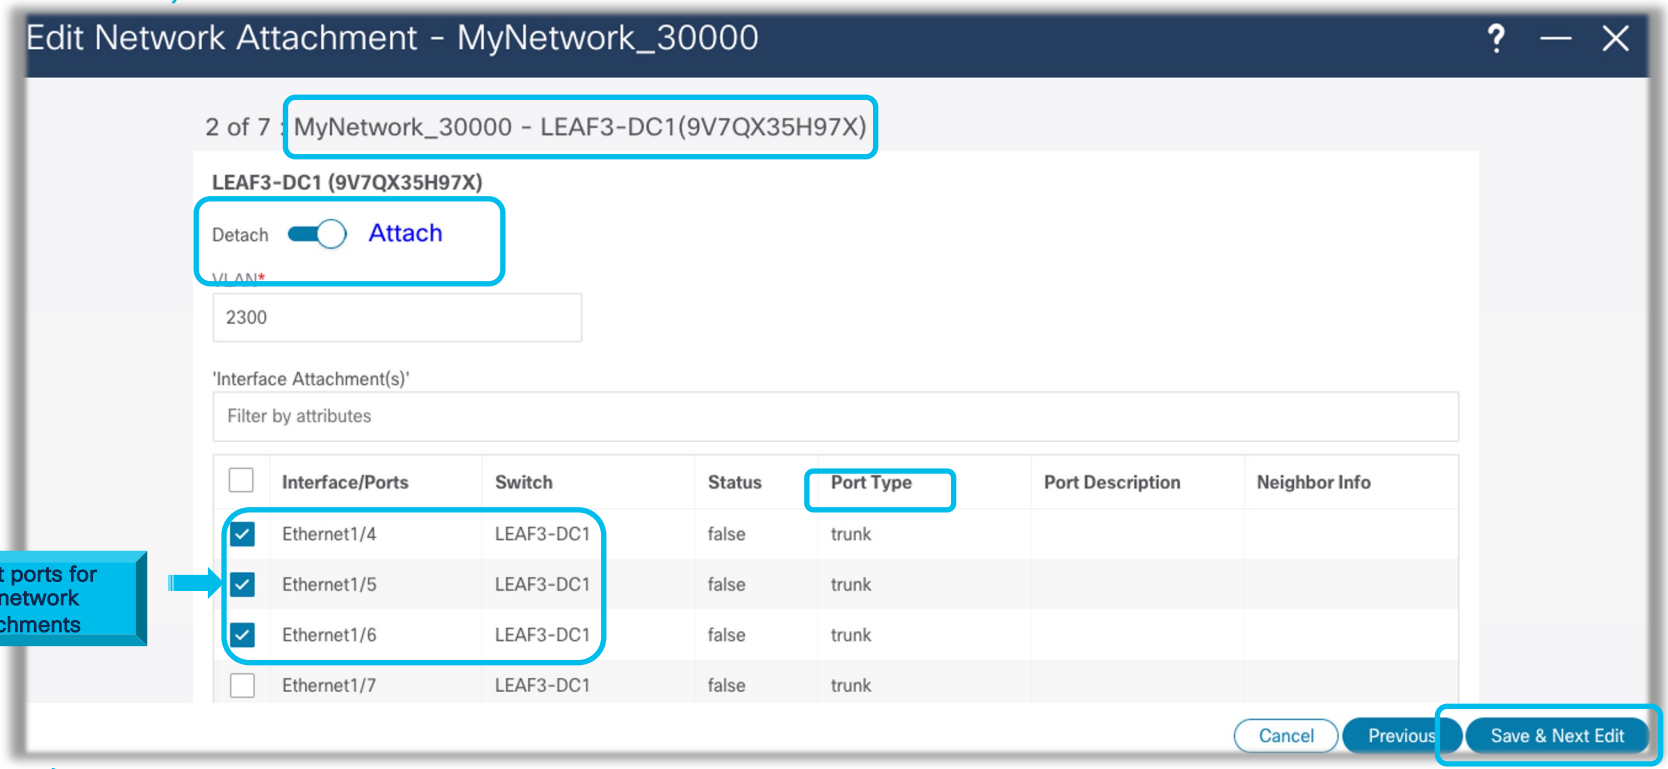

Select ports for the network attachments

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Step 4 -> Preview and Deploy VRF and Network

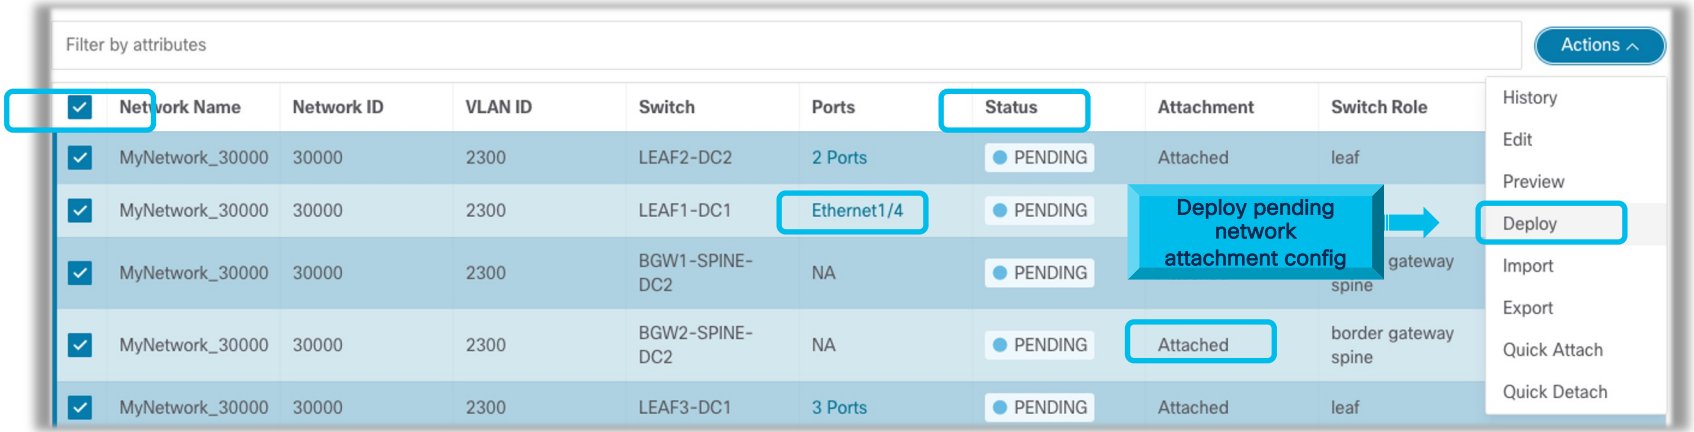

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF and Network Configs

Pending Config - VXLAN_MULTI-SITE - BGW1-DC1

L3VNI VRF Configs

```
vlan 2000
  vn-segment 50000
vrf context myvrf_50000
  vni 50000
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
exit
interface Vlan2000
  vrf member myvrf_50000
  ip forward
  ipv6 address use-link-local-only
  no ip redirects
  no ipv6 redirects
  mtu 9216
  no shutdown
```

VXLAN Tenant

```
router bgp 65001
  vrf myvrf_50000
    address-family ipv4 unicast
      advertise l2vpn evpn
      redistribute direct route-map fabric-rmap-redist-subnet
      maximum-paths ibgp 2
      exit
    address-family ipv6 unicast
      advertise l2vpn evpn
      redistribute direct route-map fabric-rmap-redist-subnet
      maximum-paths ibgp 2
```

NVE Tunnel Configs

```
configure terminal
interface nve1
  member vni 30000
    multisite ingress-replication
    mcast-group 239.1.1.1
  member vni 50000 associate-vrf
```

L2VNI Network Configs

```
vlan 2300
  vn-segment 30000
evpn
  vni 30000 l2
    rd auto
    route-target import auto
    route-target export auto
```


© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF and Network Configs

Pending Config - VXLAN_MULTI-SITE - LEAF3-DC1

L3VNI VRF Configs

```
vlan 2000
  vn-segment 50000
vrf context myvrf_50000
  vni 50000
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
exit
interface Vlan2000
  vrf member myvrf_50000
  ip forward
  ipv6 address use-link-local-only
  no ip redirects
  no ipv6 redirects
  mtu 9216
  no shutdown
```

VXLAN Tenant

```
router bgp 65001
  vrf myvrf_50000
    address-family ipv4 unicast
      advertise l2vpn evpn
      redistribute direct route-map fabric-rmap-redist-subnet
      maximum-paths ibgp 2
      exit
    address-family ipv6 unicast
      advertise l2vpn evpn
      redistribute direct route-map fabric-rmap-redist-subnet
      maximum-paths ibgp 2
```

NVE Tunnel Configs

```
configure terminal
interface nve1
  member vni 30000
    mcast-group 239.1.1.1
  member vni 50000 associate-vrf
```

L3 SVI Configs

```
vlan 2300
  vn-segment 30000
interface Vlan2300
  vrf member myvrf_50000
  ip address 10.10.10.1/24 tag 12345
  fabric forwarding mode anycast-gateway
  no shutdown
exit
```

L2VNI Network Configs

```
evpn
  vni 30000 l2
    rd auto
    route-target import auto
    route-target export auto
configure terminal
interface ethernet1/4
  switchport trunk allowed vlan add 2300
interface ethernet1/5
  switchport trunk allowed vlan add 2300
interface ethernet1/6
  switchport trunk allowed vlan add 2300
```


© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF and Network Deployment Status

In VXLAN EVPN Multi-Site Fabric

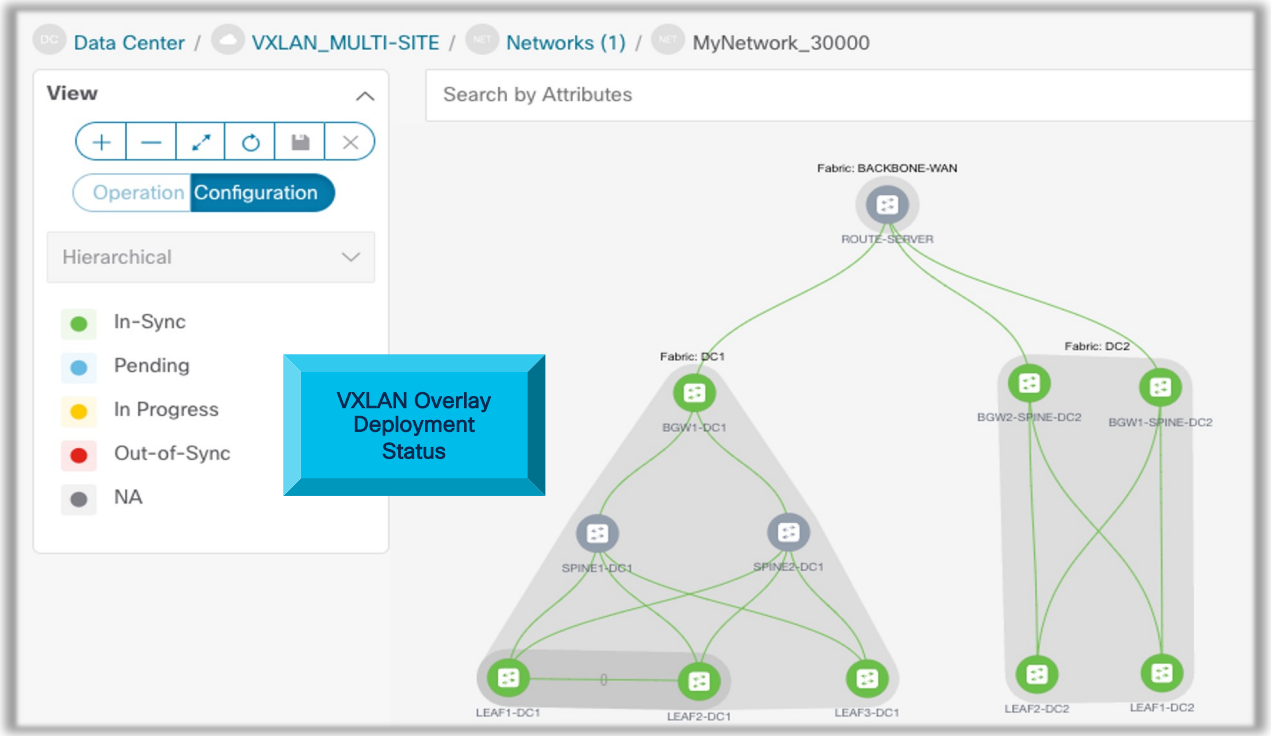

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VXLAN EVPN Brownfield Migration

Already using VXLAN?
Want to use NDFC for managing your fabric?

NDFC fully supports Brownfield	Non-disruptive import of existing VXLAN EVPN deployments	Learns topology, all configuration, associated resources, IP subnets, VNIs, VLANs, etc.	Start Managing fabric as if provisioned from NDFC

Step 1
Create

Define fabric settings (Underlay, Overlay) - **Match** AS#, Replication Mode, IGP, etc.

Step 2
Discover

Import switches with **Preserve Config**

Define switch Roles (Border, Leaf, Spine, etc)

Step 3
Recalculate and Deploy

Sanity checks for mis-config and Normalizes configuration to best practices

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Manage and Deploy External IP Handoff with NDFC

cisco Live!

External Handoff Placement

Border Role compared to Leaf Role

- **Leaf** VTEP (**East-West** traffic)
  - Connectivity to endpoints, the first-hop routing
  - Server-to-server traffic
- **Border** Leaf VTEP (**North-South** traffic)
  - Transit node
  - Traffic destined from within the VXLAN EVPN Fabric to outside or vice versa
  - Optionally it can have directly attached endpoints
- Separating Roles simplifies operations

V V = VTEP

R = RR/RP

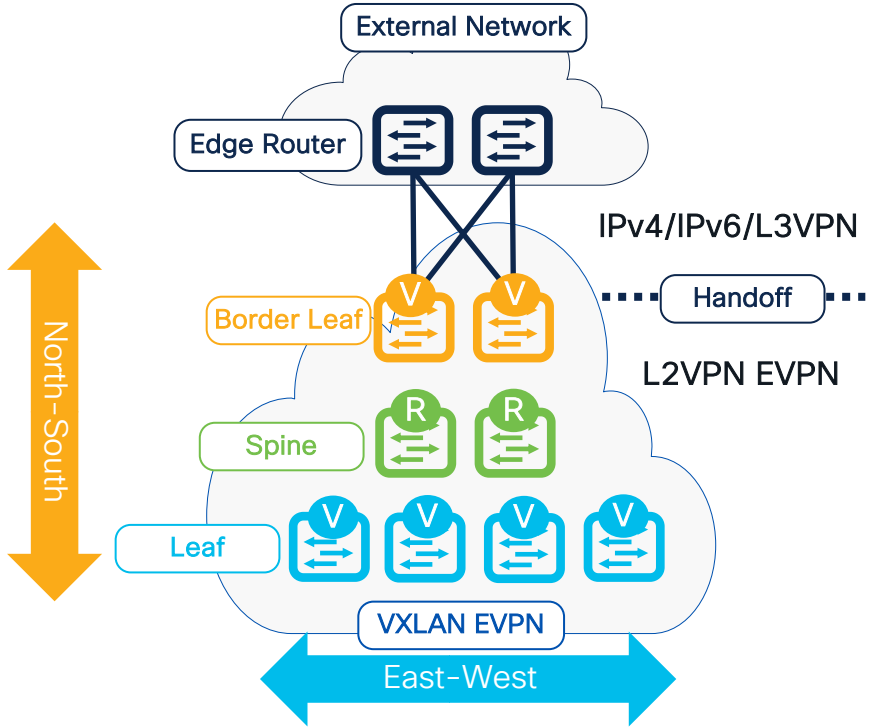

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Border Node Design
VXLAN EVPN Fabric

Border Node Placement Option

Flexible Design option-1

Border as Leaf nodes (**Border Leaf**)

- Capacity planning only for north-south traffic flows
- Clean role separation and uniform reachability from the entire fabric are the major advantages
- Support for Inter-AS option A (VRF-LITE) and seamless VXLAN-MPLS gateway (Border-PE)

Border Leaf hosts:

VTEP Mainly for

- North-South

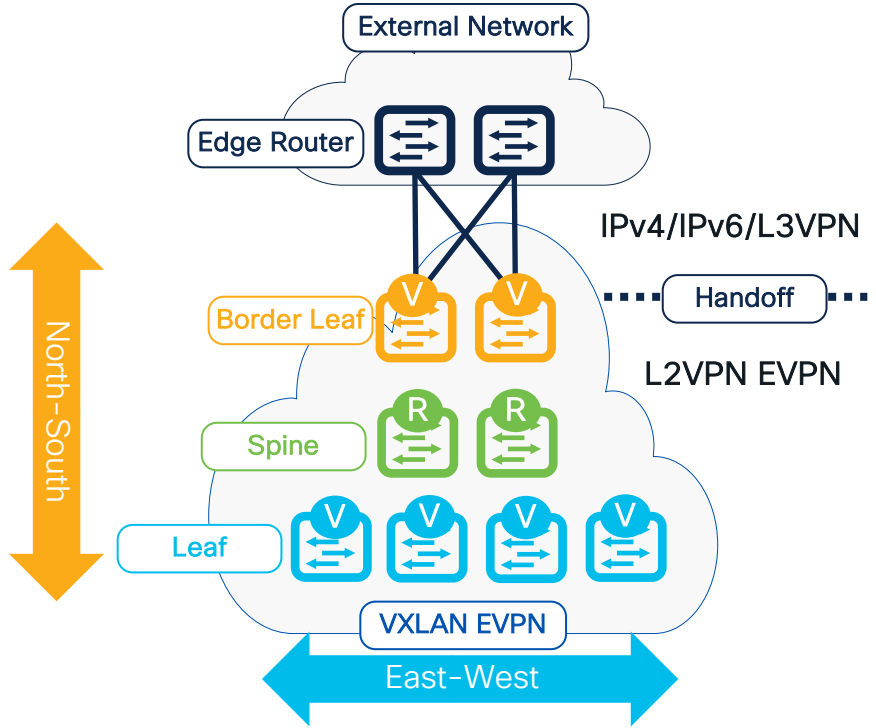

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Border Node Placement Option

Flexible Design option-2

Border as Spine nodes (**Border Spine**)

- Leaf endpoints are one-hop away from External Networks
- Allowance for a scale-out approach to add spines and provide more bandwidth and resiliency
- Support for Inter-AS option A (VRF-LITE) and seamless VXLAN-MPLS gateway (Border-PE)
- Extra functional dependency (Border + Spine)
- Capacity planning needs to accommodate all flows

Border Spine hosts:
VTEP for:
- North-South
- East-West

- Route Reflector (RR) (iBGP EVPN)
- Rendezvous Point (RP) (Multicast Underlay BUM)

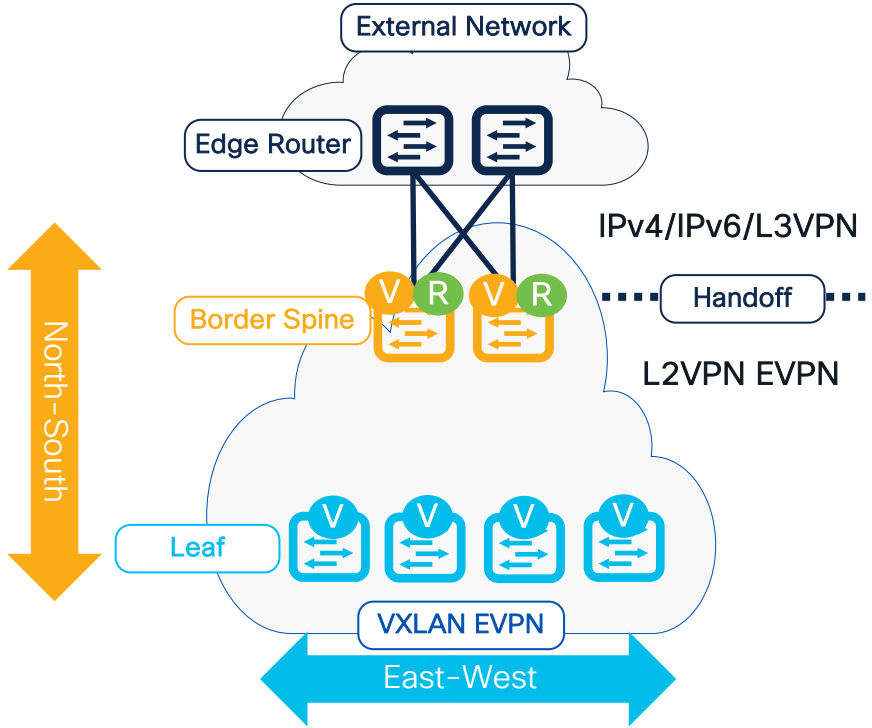

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Who can be the External Edge Router?

Edge Router Placement Option

Nexus and Non-Nexus support

TIP

Supported Hardware and Software might vary depending on NDFC version
Check compatibility matrix 12.1.3b

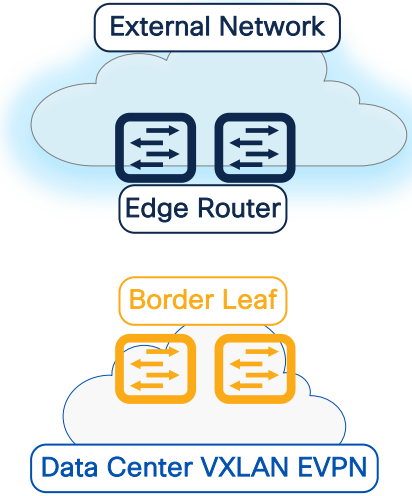

Managed NX-OS as Edge Router
- Nexus 9000
- Nexus 7000

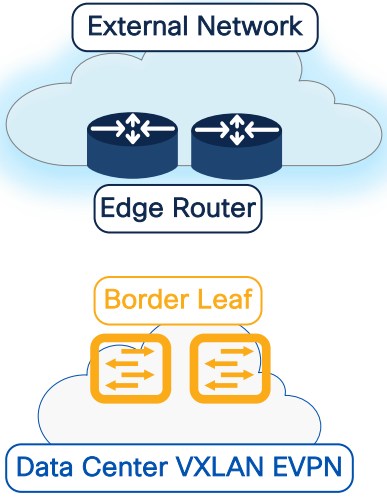

Managed IOS-XR as Edge Router
- Cisco 8000
- ASR 9000
- NCS 5500

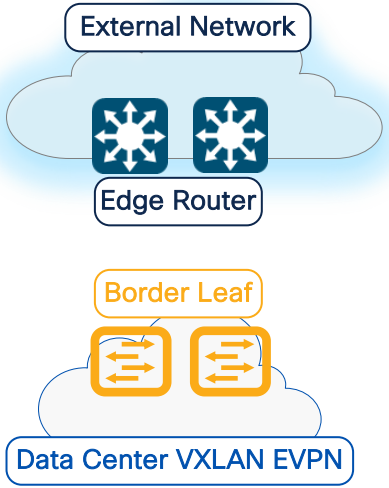

Managed IOS-XE as Edge Router
- Catalyst 9000
- Catalyst 8000
- ASR 1000
- CSR 1000

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Border to External
Peering Type

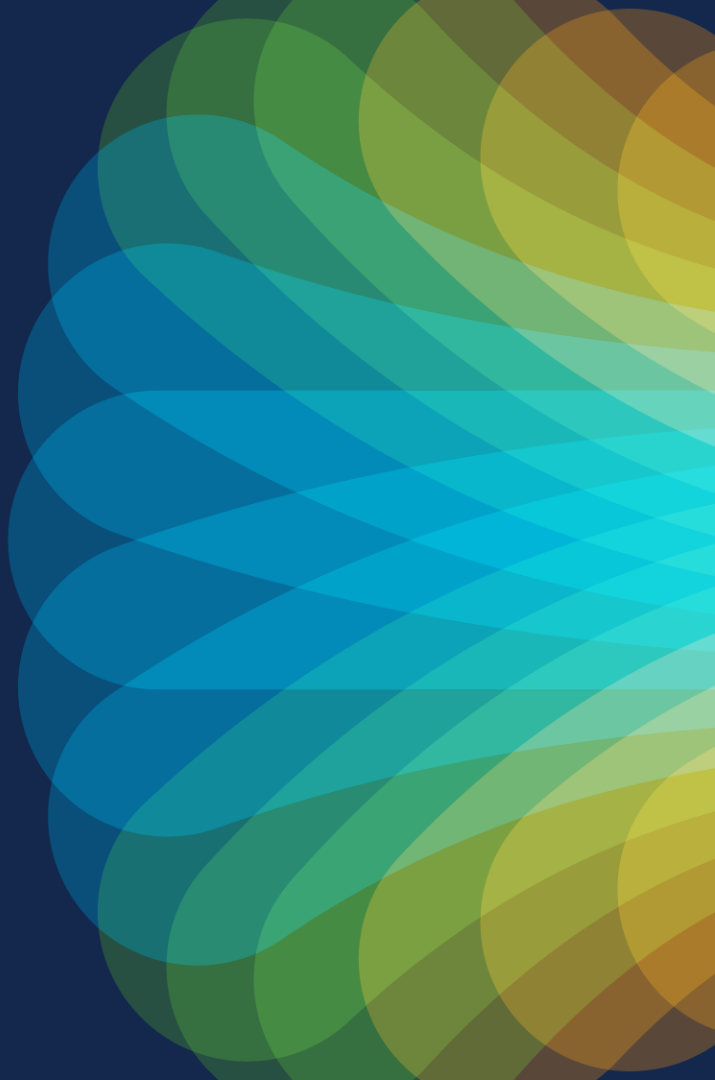

External Layer-3 connectivity options

Inter-AS Option A

- Clear separation of Autonomous Systems
- Simple, Straight forward, and Commonly used
- No need for redistribution
- Easy and Flexible BGP route-filtering mechanisms
- BGP natural loop avoidance
- Structured handoff between the VXLAN BGP EVPN fabric and the external routing domain (Backbone, WAN, Campus, etc.)
- Not ideal for High scale VRF handoff deployment

Peering Type = Sub-interfaces on physical routed (or L3 Port-channel) interfaces

- Sub-interface with dot1q tag to mark the traffic to a specific VRF
- Sub-interface used for eBGP peering and as next-hop
- Per VRF, Per Sub-interface eBGP session

VPNv4/VPNv6

IPv4/IPv6

EVPN

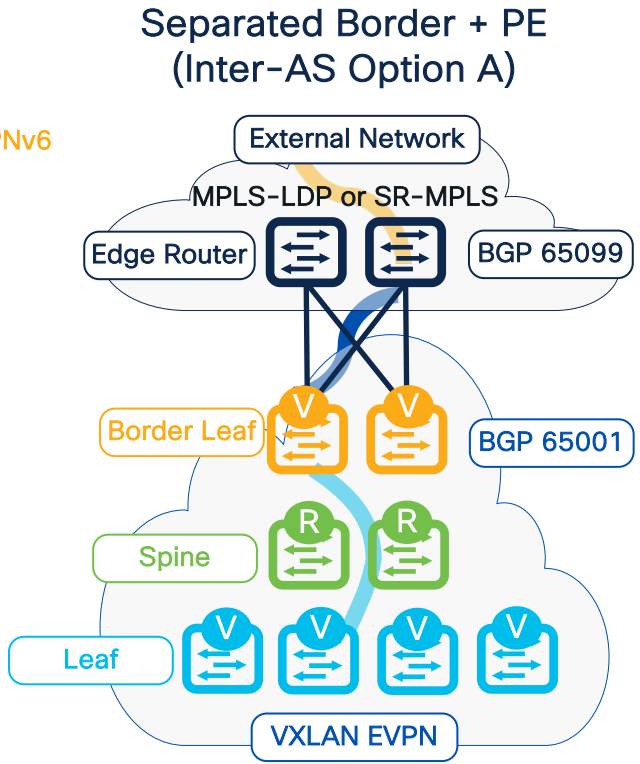

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

External Layer-3 connectivity options

Seamless Protocol Gateway Model (Border-PE)

- Combines two different encapsulations and Address Family, using a “single-box (Border-PE)” instead of a “two-box (CE-PE)” model
- VXLAN VTEP Border nodes also becomes a MPLS L3VPN Provider Edge (PE), resulting in a role called Border-PE
- Best suited for high scale VRF deployment
- Saves CAPEX and OPEX
- Seamless stitching between L2VPN EVPN and VPNv4/v6 Address Family
- BGP route-filtering mechanisms available
- Specific Hardware support
  - MPLS LDP: Nexus 3600-R, Nexus 9500-R
  - SR MPLS: Nexus 9300 FX2/FX3/GX/GX2, Nexus 9500-R

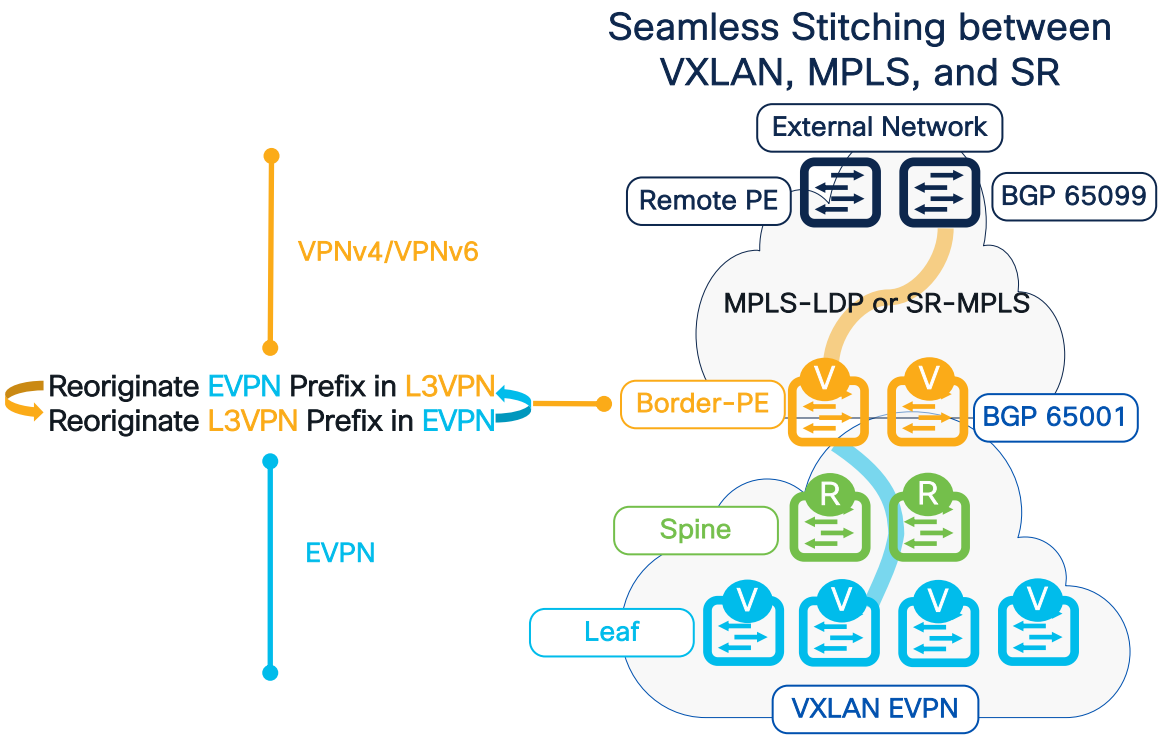

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Provision External IP Handoff with NDFC

cisco Live!

Enabling VRF-LITE Automated

Border Leaf to Nexus Edge Router

Step 1

Create and Import

Define VXLAN EVPN and External Fabric. Set respective roles (e.g., Border, BGW, Edge)

Uncheck Fabric monitor mode in External Fabric if the Edge router is in managed mode

Step 2

IFC prototypes

Define IFC Type: "Back2Back&ToExternal"

Check: Auto-Deploy both

(Optional) Auto Allocation of Unique IP on VRF Extension over VRF LITE IFC

Recalculate config and deploy

Step 3

Deploy Symmetric VRF-Lite Extension

Define Individual VRF extension on the Border leaf

NDFC will add Sub-int and eBGP peering on Border leaf and neighbor Edge router

Deploy on Data Center VXLAN EVPN Fabric and External Network Fabric

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Topology and IFC considerations

- IFC Automated
- Advertise Host (disabled)
- Advertise Default-Route (enabled)
- Config Static 0/0 Route (enabled)

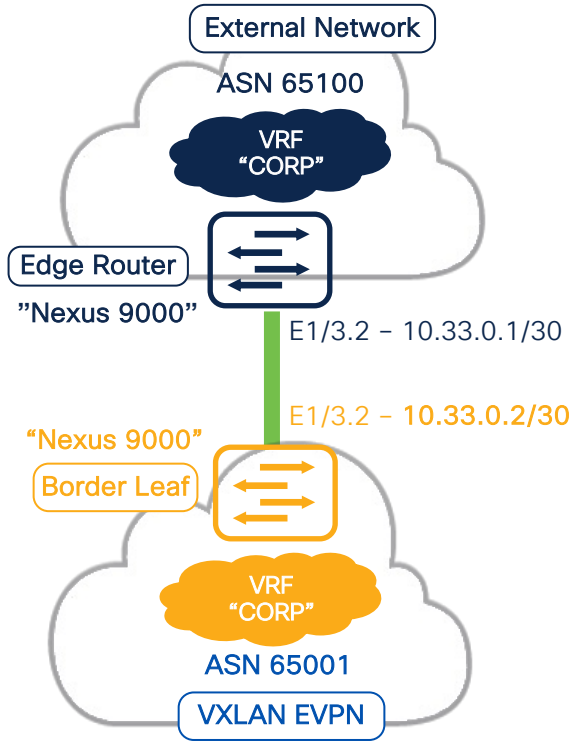

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Managing Edge Devices

part of “External Connectivity Network” Fabric

- Create Fabric, Discover Switches, Set Role, and Recalculate & Deploy

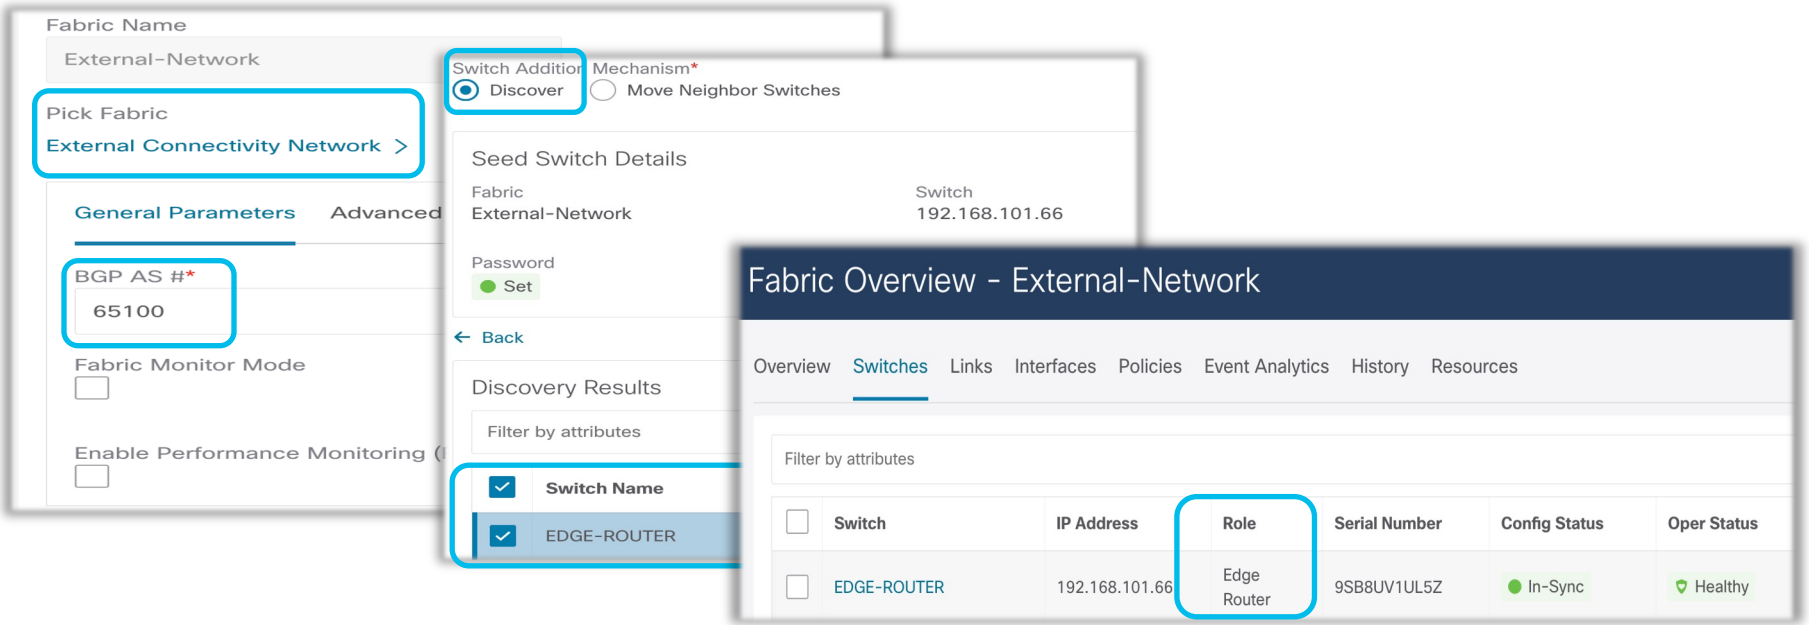

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining IFC deployment type

- Review Fabric Settings for **VRF-Lite IFC deployment type:**

  LAN > Fabrics > Select (your DC VXLAN Fabric) > Actions > Edit Fabric > **Resources Tab**

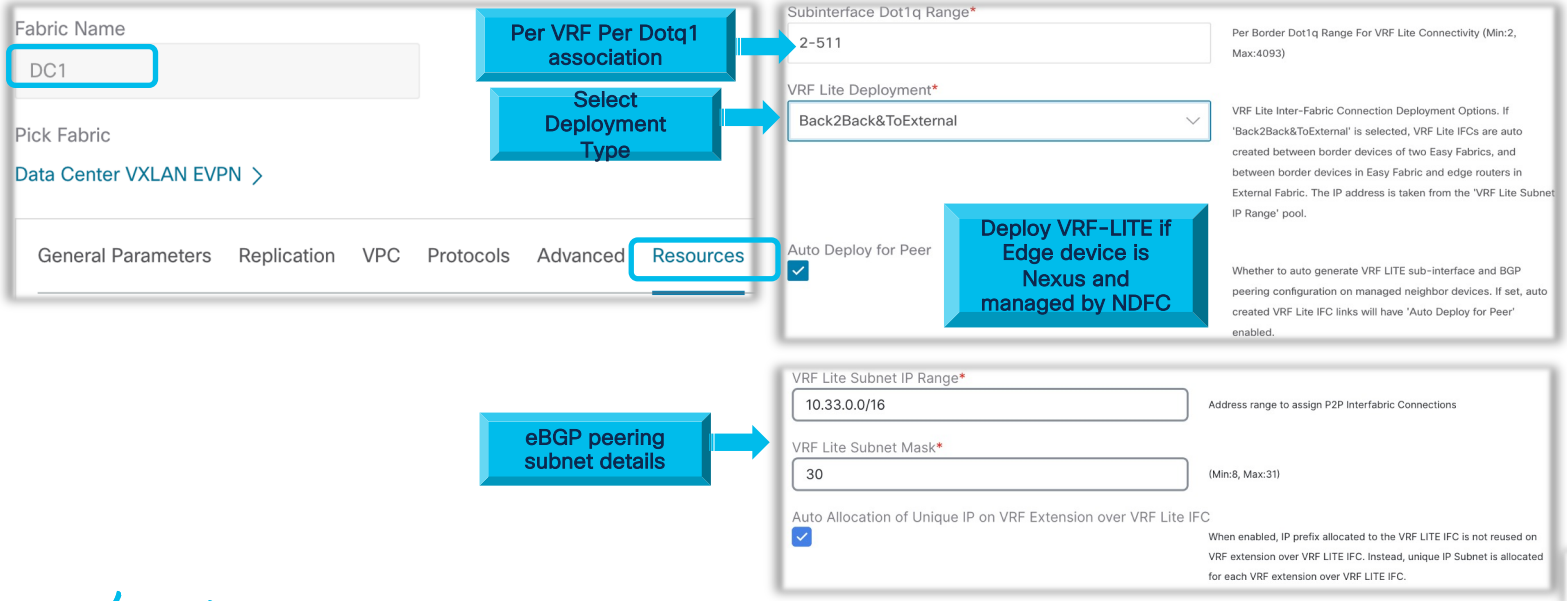

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining IFC Link on physical interface

- IFC link has been defined
  - Policy should be **ext_fabric_setup**
    - IPs auto selected from **VRF-Lite Subnet IP Range**

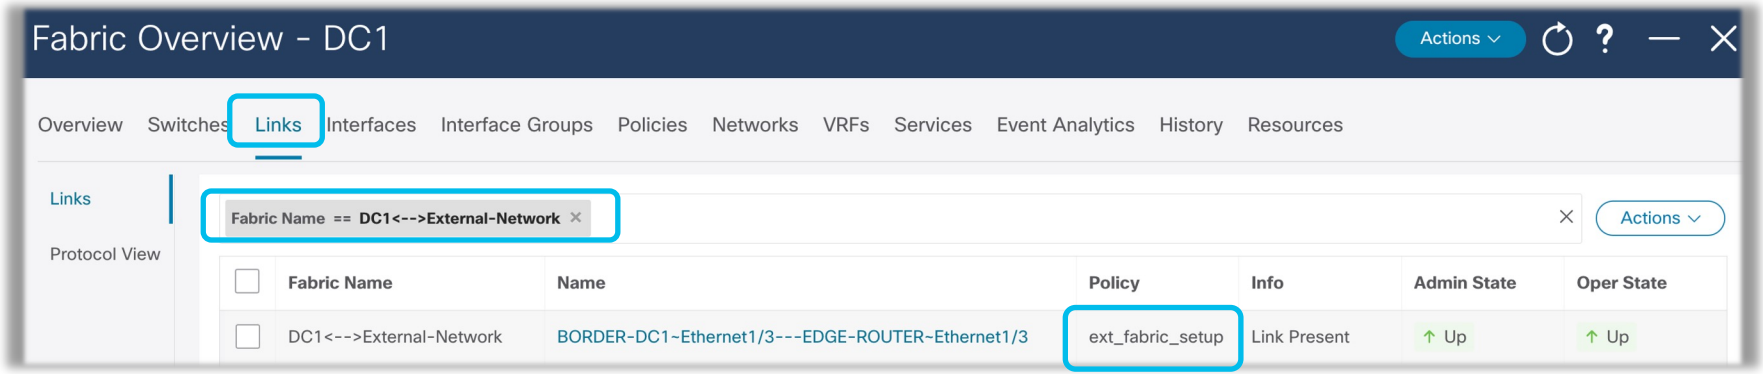

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining VRF extensions on Border

- Verify DC VXLAN Fabric VRFs were created and customize if needed
  - LAN > Fabrics > **Double click** (your DC VXLAN Fabric) > **VRFs**

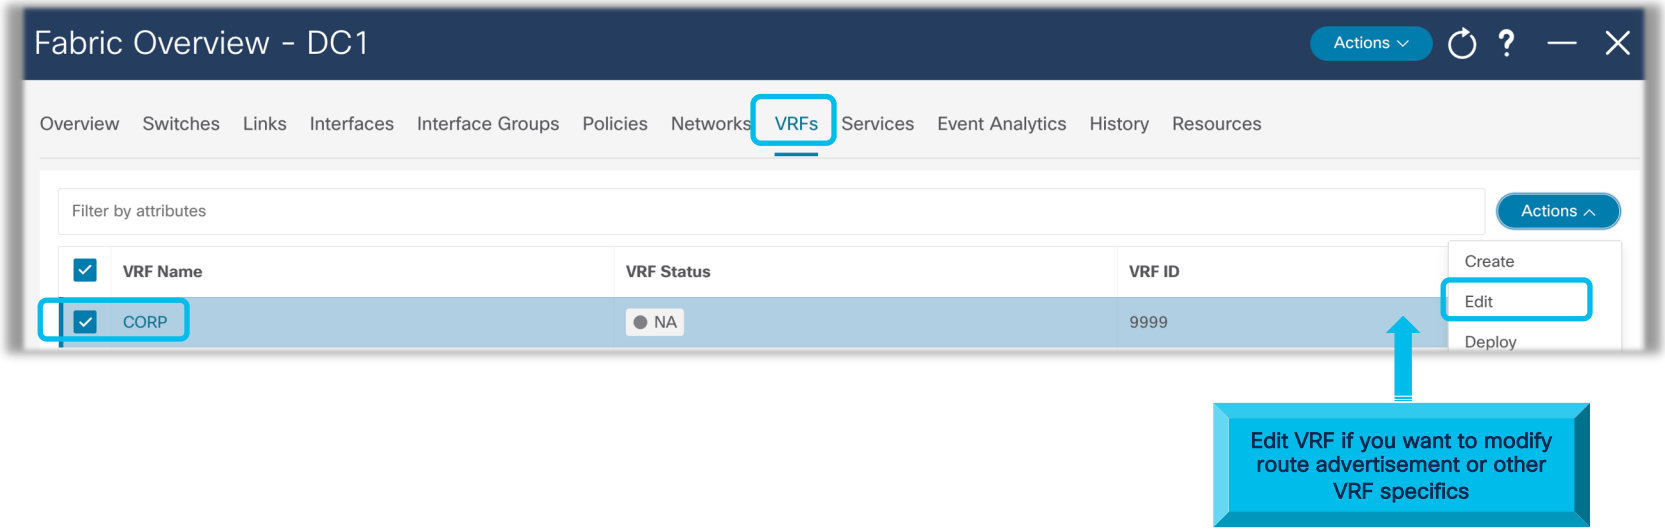

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining VRF extensions on Border

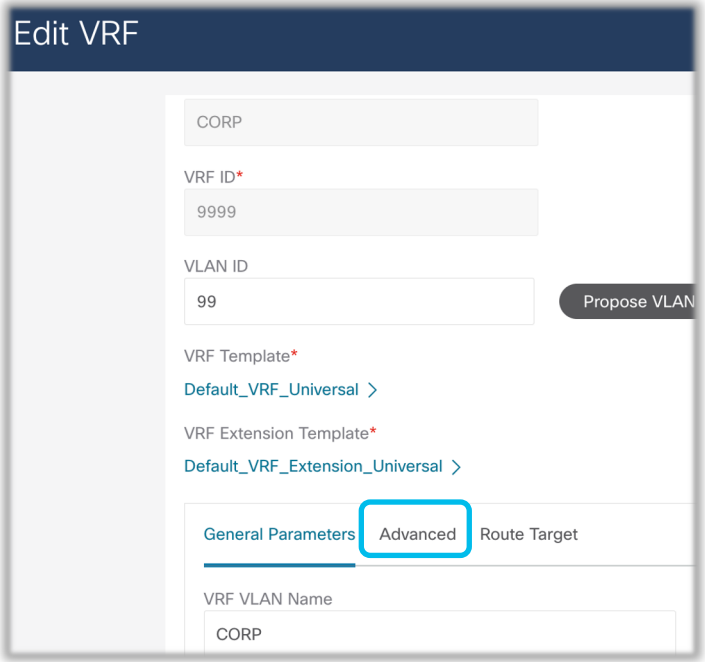

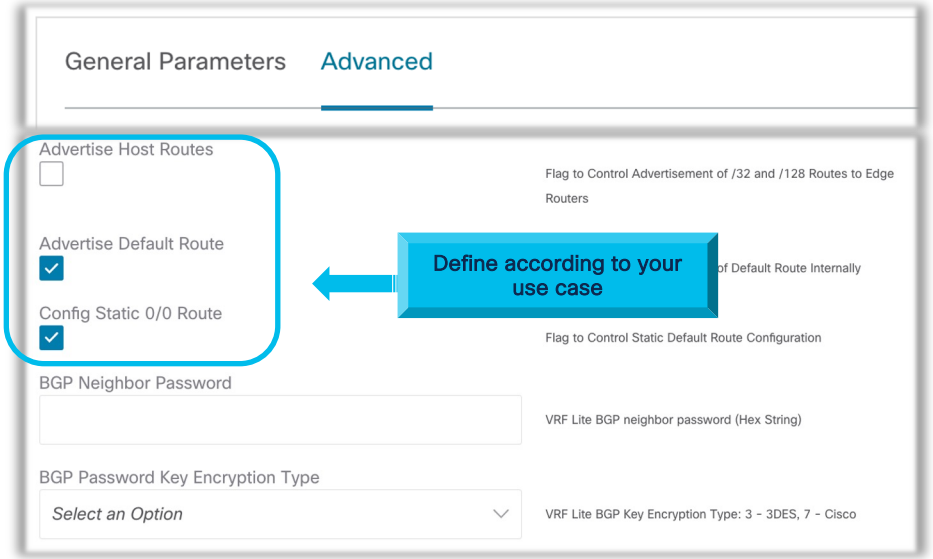

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining VRF extensions on Border

- **Attach** VRF to Border and extend through **VRF-Lite**

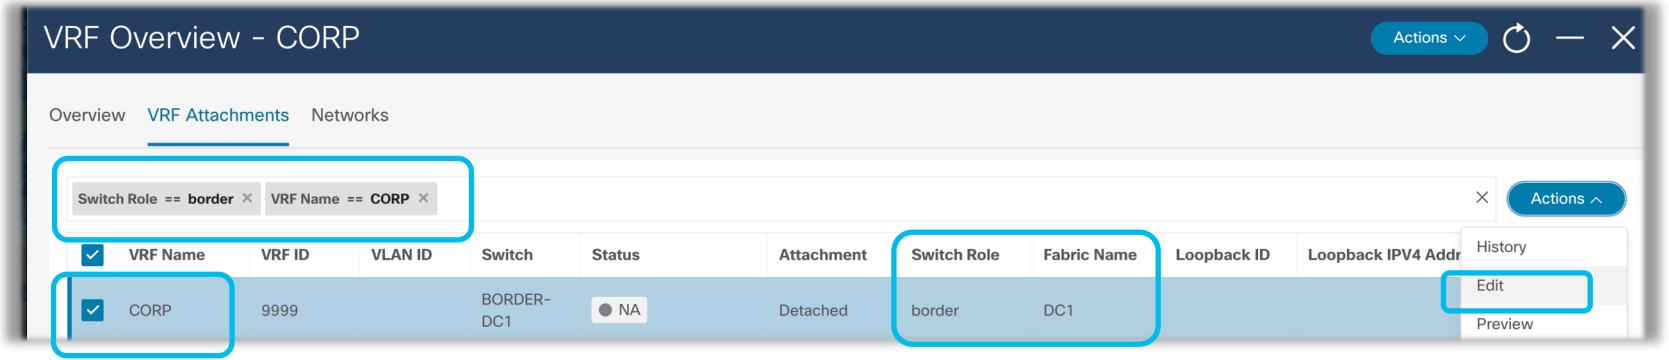

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining VRF extensions on Border

- **Attach** VRF to Border and extend through **VRF-Lite**

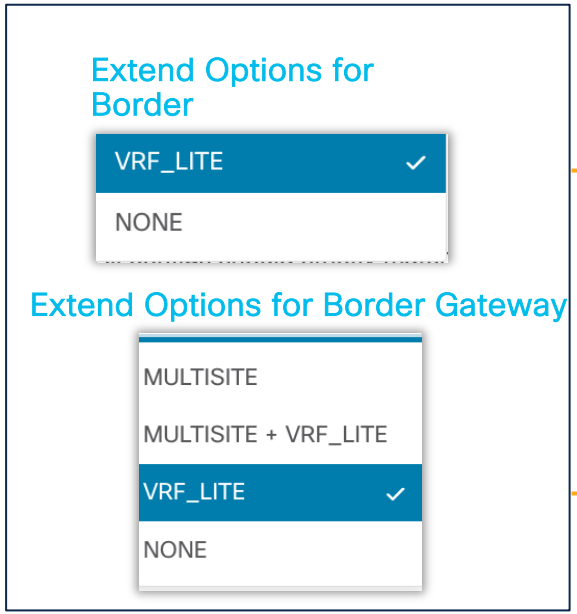

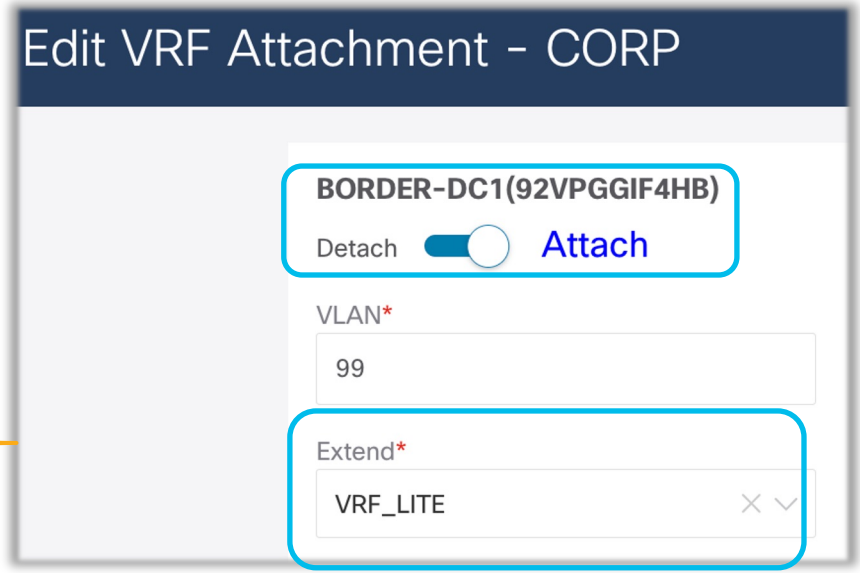

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Defining VRF extensions on Border

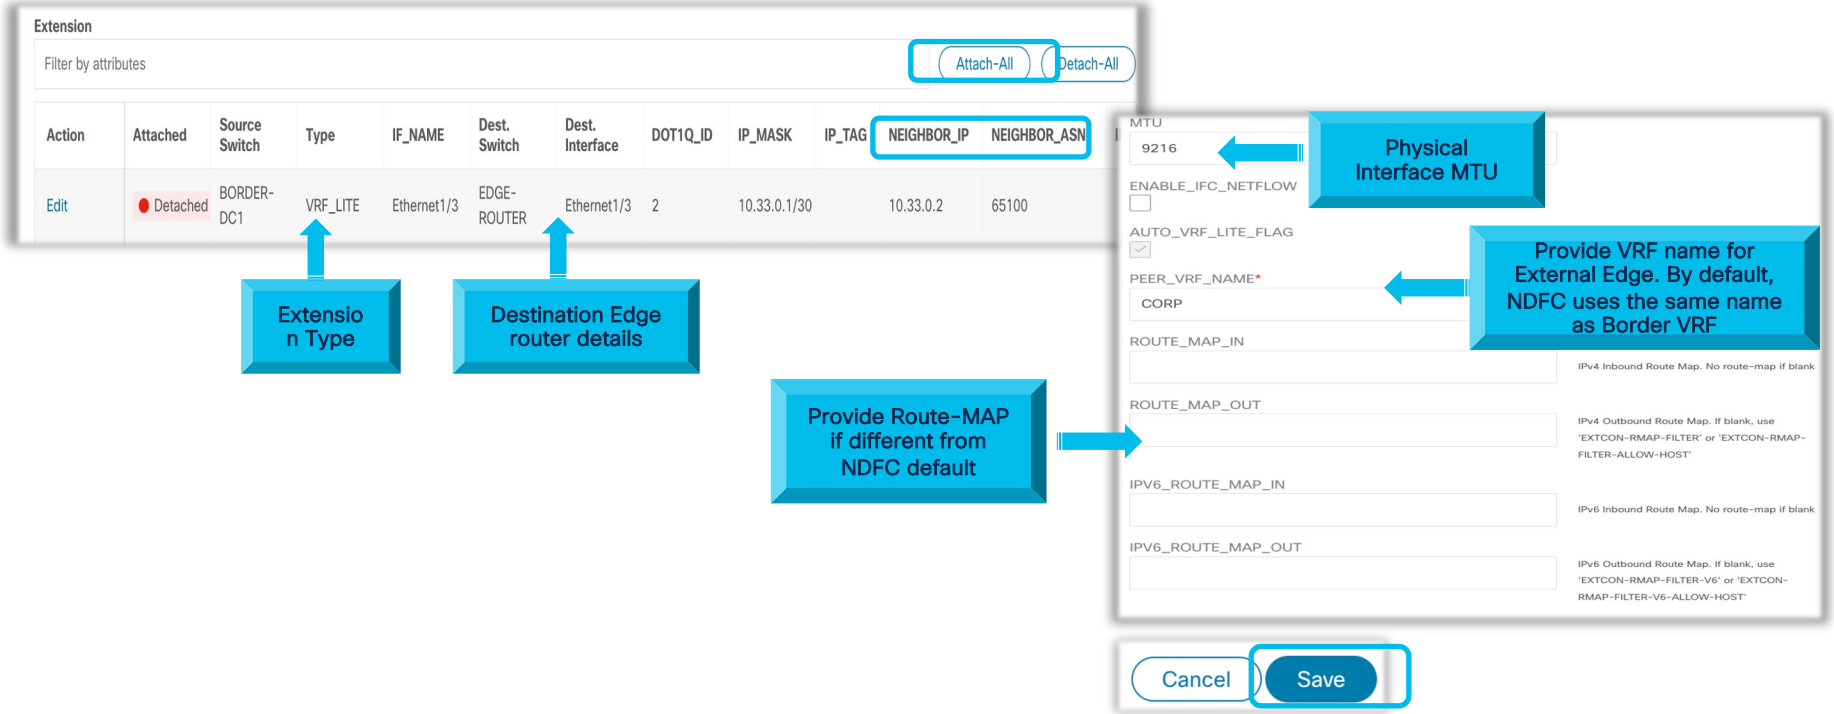

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Preview and Deploy VRF extensions on Border

Pending Config - DC1 - BORDER-DC1

Pending Config | Side-by-Side Comparison

```
vlan 99
  name CORP
  vn-segment 9999
vrf context corp
  description CORP
  vni 9999
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  ip route 0.0.0.0/0 10.33.0.2
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
exit
interface Vlan99
  description CORP
  vrf member corp
  ip forward
  ipv6 address use-link-local-only
  no ip redirects
  no ipv6 redirects
  mtu 9216
  no shutdown
```

L3VNI VRF Configs

Default route static towards External Edge

L3 Tenant CORP

```
router bgp 65001
 vrf corp
   address-family ipv4 unicast
     advertise l2vpn evpn
     redistribute direct route-map fabric-rmap-redist-subnet
     maximum-paths ibgp 2
     network 0.0.0.0/0
     exit
   address-family ipv6 unicast
     advertise l2vpn evpn
     redistribute direct route-map fabric-rmap-redist-subnet
     maximum-paths ibgp 2
     exit
   neighbor 10.33.0.2
     remote-as 65100
     address-family ipv4 unicast
       send-community both
       route-map extcon-rmap-filter out
```

Advertisement of Default route towards site internal VTEP

eBGP External Edge neighbor

```
interface ethernet1/3.2
  encapsulation dot1q 2
  mtu 9216
  vrf member corp
  ip address 10.33.0.1/30
  no shutdown
interface nve1
  source-interface loopback1
  host-reachability protocol bgp
  no shutdown
  member vni 9999 associate-vrf
```

Border Physical Interface Config

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

VRF-LITE: Border to Nexus Edge

Preview and Deploy VRF extensions on External Edge

- Once configurations are deployed on Border Leaf (DC VXLAN EVPN), navigate to **External Network Fabric** and **perform Recalculate and Deploy**

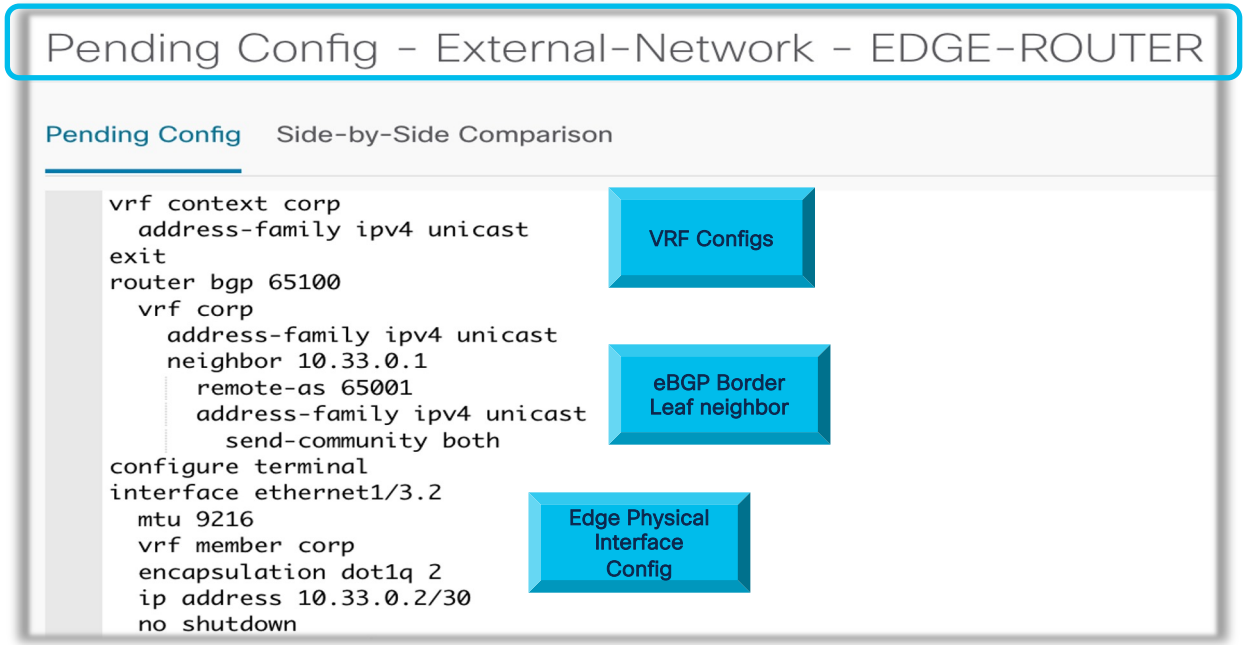

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Verification and Validation with NDFC

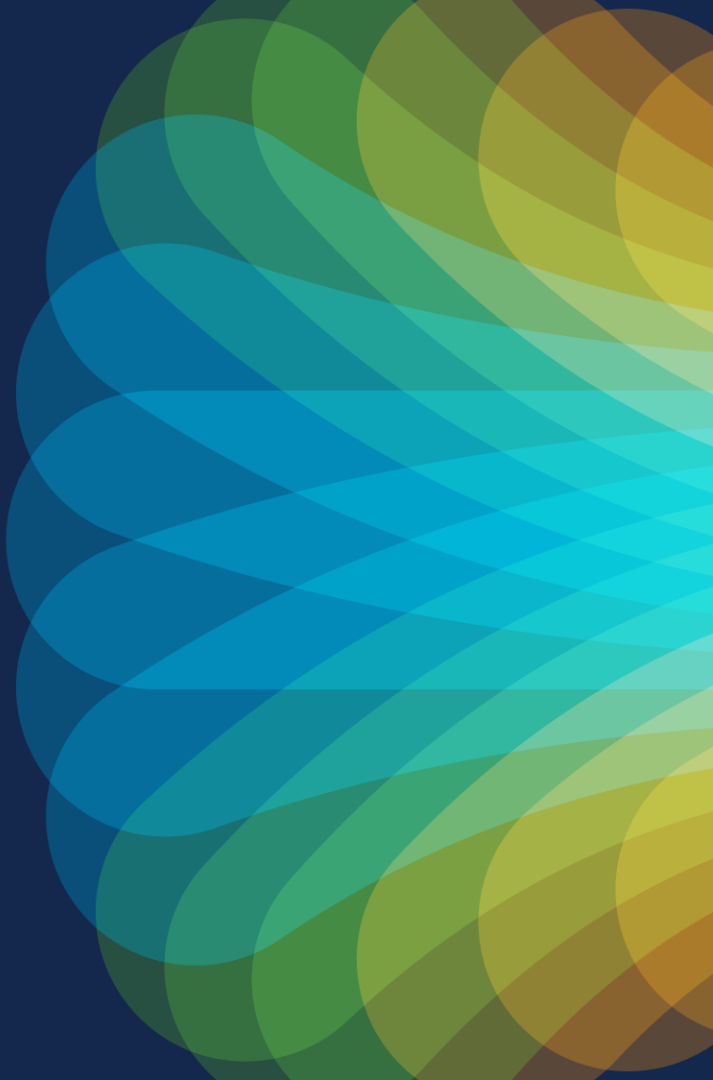

Verification through NDFC

Keeping you away from CLI

Step 1
Verify Network and VRF attachments

Configuration Execution Status:

Network Status **Deployed**

VRF Status **Deployed**

Step 2
Deployment History

Configuration Execution Status:

Verify Deployment History Status **Success** for Underlay, Overlay, Interfaces, and more

Step 3
Show commands

Service / features status

(CLI through NDFC)

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Attachment deployment status

Job execution perspective

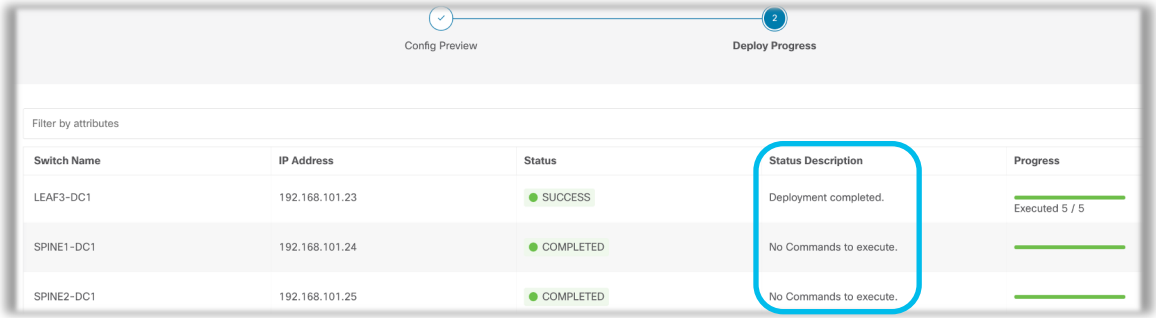

Success or Failure deployment details

Multi-Stage Preview and Deployment

Switch Name	IP Address	Status	Status Description	Progress
Edge-router	192.168.101.99	FAILED	Deployment failed. Check deployment history for more information.	Executed 1 / 12
Edge-Catalyst	192.168.101.101	COMPLETED	No Commands to execute.	

Role	VRF Status	Status Description	Progress
border gateway	Deployment In-Pr	Adding diff to deployment queue	

Role	VRF Status	Status Description	Progress
border gateway	In-Sync	Config compliance sync completed	

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Deployment History Tool

Commands execution perspective

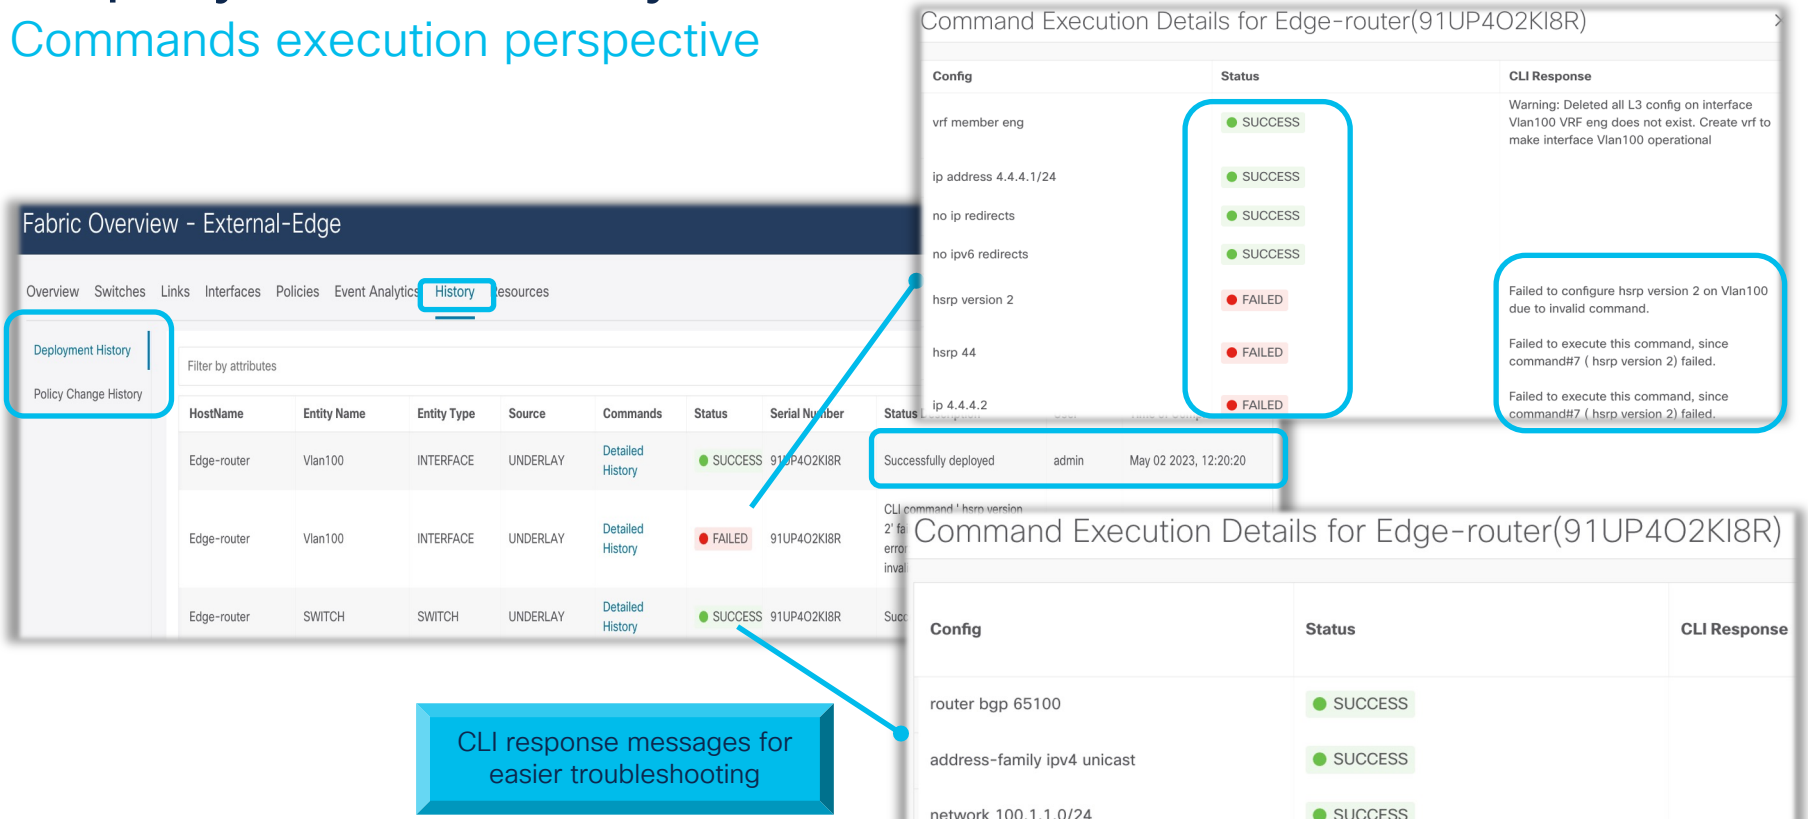

CLI response messages for easier troubleshooting

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Show Commands Tool

Switch Config perspective

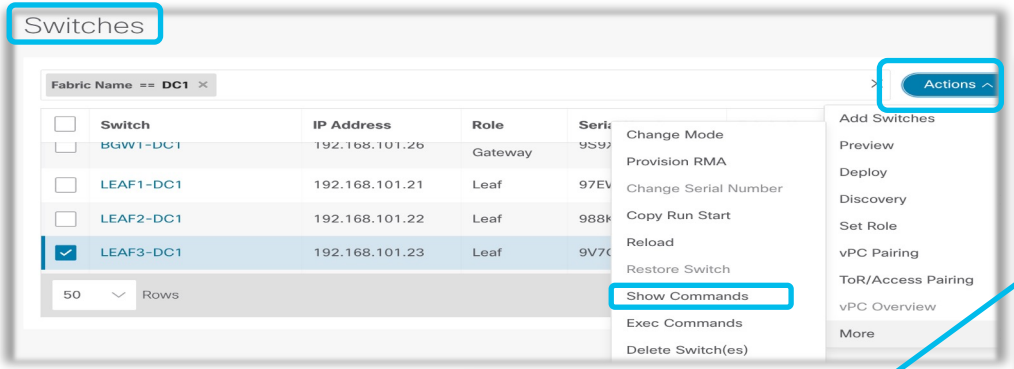

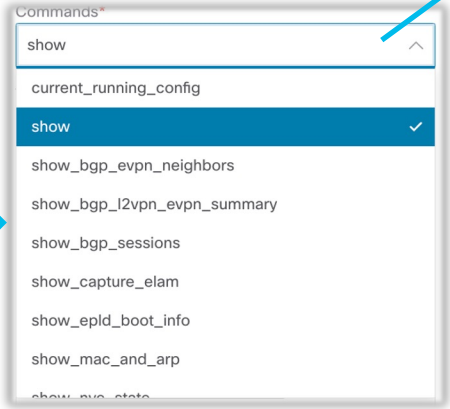

NDFC pre-built commands, user commands, or custom-built show templates

```
#show ip ospf nei
OSPF Process ID UNDERLAY VRF default
 Total number of neighbors: 2
 Neighbor ID     Pri State            Up Time  Address         Interface
 10.2.0.6          1 FULL/ -          2d22h    10.4.0.29       Eth1/2
 10.2.0.1          1 FULL/ -          2d22h    10.4.0.21       Eth1/3

#show bgp l2vpn evpn su
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 10.2.0.3, local AS number 65001
BGP table version is 29, L2VPN EVPN config peers 2, capable peers 2
9 network entries and 16 paths using 2436 bytes of memory
BGP attribute entries [16/2752], BGP AS path entries [1/10]
BGP community entries [0/0], BGP clusterlist entries [6/24]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.2.0.1        4 65001    4252    4240       29    0    0    2d22h 5
10.2.0.6        4 65001    4252    4239       29    0    0    2d22h 5

#show int nve1
nve1 is up
admin state is up,  Hardware: NVE
  MTU 9216 bytes
  Encapsulation VXLAN
```


© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Easily license your NDFC

Software included with switch subscription licenses

Tier-based

NXOS Essentials, Advantage, Premier

OR

DCN Essentials, Advantage, Premier

Install base (NX-OS) | New purchase (NX-OS)

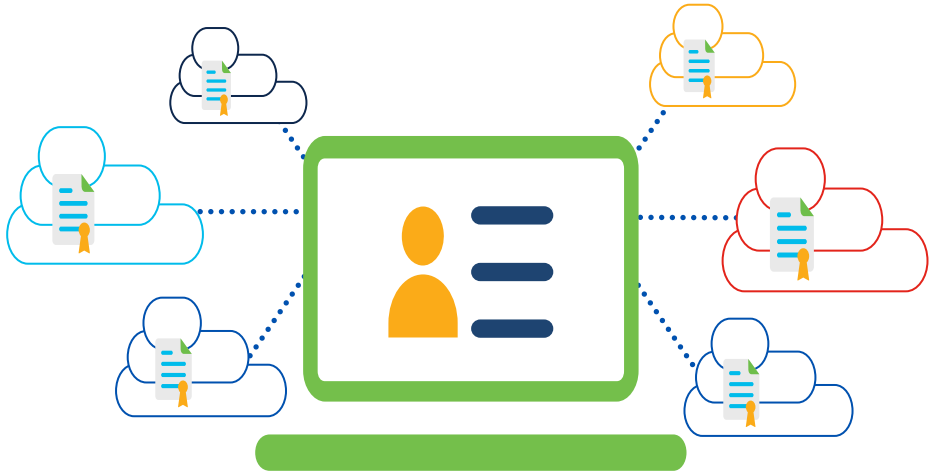

- Increased ease of use and efficient consumption
- Less restrictions with flexible reporting
- Reduced overall licensing friction for users

NDFC 12.0 is transitioning to a Smart Licensing only release

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

APPENDIX

White Paper and Document References

- VXLAN BGP EVPN
  - https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/white-paper-listing.html
- NDFC
  - https://www.cisco.com/c/en/us/products/cloud-systems-management/prime-data-center-network-manager/white-paper-listing.html

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Conclusion

Key points to remember

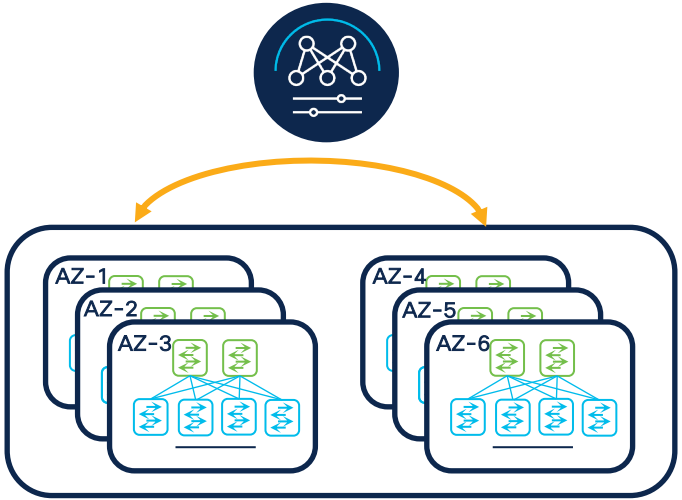

- VXLAN EVPN Multi-Site maintains clear change and fault domain separation to deploy large-scale and highly available DC architectures.
- NDFC provides simplified mechanism to extend and provide end-to-end network and policy consistency across regions, all with a single point of orchestration.
- NDFC provides a single plane of glass solution to automate and manage Nexus and Non-Nexus devices

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

cisco Live!

Did you know?

You can have a
one-on-one session with
a technical expert!

Visit Meet the Expert in The HUB
to meet, greet, whiteboard & gain
insights about your unique questions
with the best of the best.

Meet the Expert Opening Hours:

Tuesday	**3:00pm – 7:00pm**
Wednesday	**11:15am – 7:00pm**
Thursday	**9:30am – 4:00pm**
Friday	**10:30am – 1:30pm**

Session Surveys

We would love to know your feedback on this session!

- Complete a minimum of four session surveys and the overall event surveys to claim a Cisco Live T-Shirt

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Expert meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

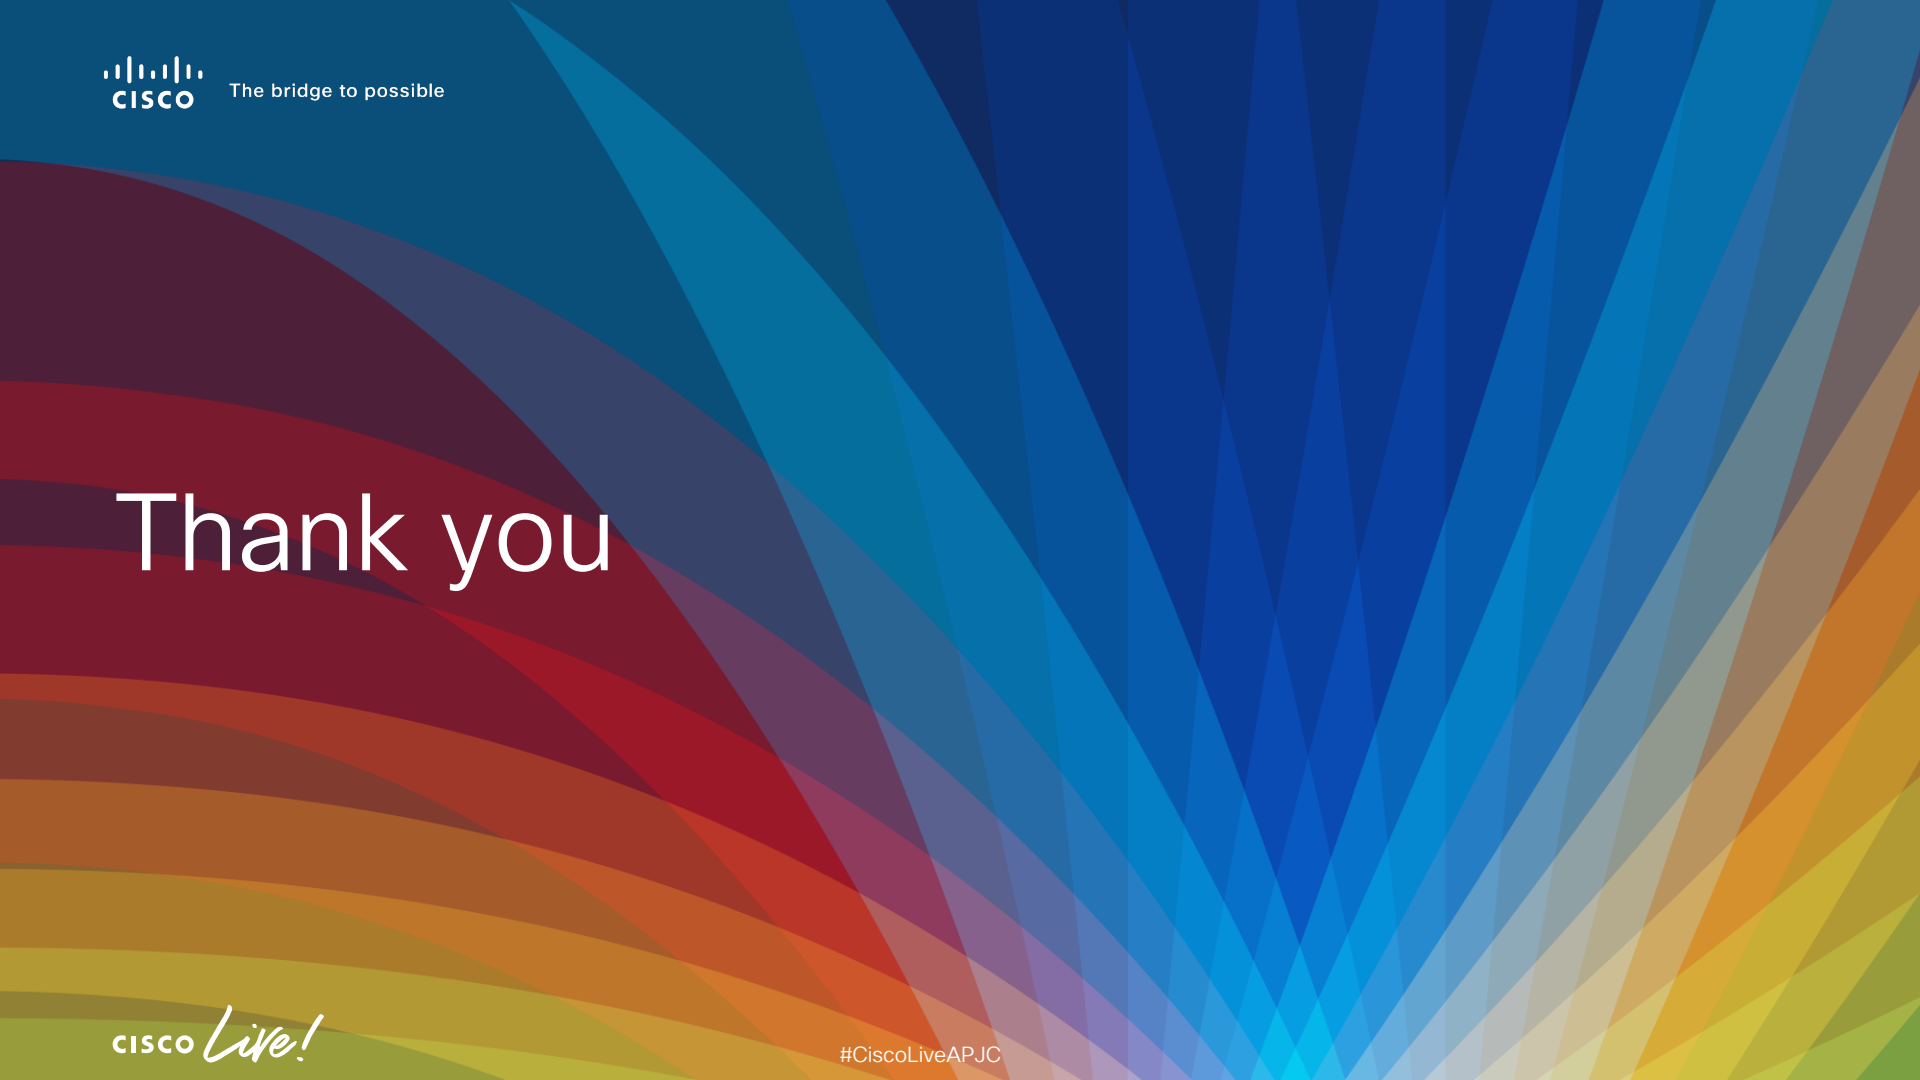
cisco
The bridge to possible
Thank you
cisco Live!
#CiscoLiveAPJC

cisco Live!

Let's go

#CiscoLiveAPJC