

Design, Automate, and Simplifying VXLAN EVPN Multi-Site Architectures with Nexus Dashboard Federation

cisco Live !

Parth Patel
Sr. Technical Leader, TME
Data Center Networking BU

Webex App


Questions?

Use Webex App to chat with the speaker after the session

How

- 1 Find this session in the Cisco Events 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 February 27, 2026.




Agenda

- 01 Design VXLAN Multi-Site with Nexus Dashboard**
- 02 Introduction to Nexus Dashboard Federation**
- 03 Automate VXLAN Multi-Site with Nexus Dashboard Federation**
- 04 Conclusion**

VXLAN Multi-Site Design Options

VXLAN EVPN Multi-Site Overview

Functional Components



Border Gateway

Deployment Considerations

Border Gateways main functions and use-cases:

- Packet Re-Origination (L2 and L3) 

- Inter-Site DCI (East-West)




- L3 Extension (North-South) 

- Connect L4-L7 services and EPs



- Integration with Legacy Networks (Co-existence and/or Migration)




- VXLAN L3-extension to the Public Cloud




Border Gateway

When to use What



Anycast Border Gateway

- Up to 6 BGW
 - Simple Failure Scenarios
- Any Deployments
 - No End-Point or Network Services Connectivity on BGW
- Greenfield Deployments



VPC Border Gateway

- 2 BGW with Physical Peer-Link
- Small Deployments
 - End-Point or Network Services Connectivity on BGW
- Migration Use-Cases (Brownfield)
 - Classic Ethernet / FabricPath to VXLAN EVPN

BGW Node Placement Option

Flexible Design option-1

Dedicated Border Gateway

- ✓ Flexible scale-out approach for VXLAN EVPN Multi-Site DCI
- ✓ Flexible Anycast or VPC BGW models
- ✓ Capacity planning only for DCI traffic flows
- ✓ Clean role separation and uniform reachability from the entire fabric are the major advantages


Border Gateway hosts:

VTEP for:

- East-West (DCI Packet Re-Origination L2/L3)

V V = VTEP

R = RR/RP



BGW Node Placement Option

Flexible Design option-2

BGW on Spine nodes (**Border Gateway Spine**)

- ✓ Flexible scale-out approach for VXLAN EVPN Multi-Site DCI
- ✓ Anycast BGW only
- ⚠ Extra functional dependency (BGW + Spine)
- ⚠ Capacity planning needs to accommodate all flows


Border Gateway Spine hosts:

VTEP for:

- East-West (DCI Packet Re-Origination L2/L3)
- Route Reflector (RR) (Site-Internal iBGP EVPN)
- Rendezvous Point (RP) (Site-Internal Multicast Underlay BUM)

V V = VTEP

R = RR/RP



BGW Node Placement Option

Flexible Design option-3

BGW on top of Super-Spine nodes

- ✓ Scale-out Multi-Clos Fabric to Interconnect the PODs using Super-Spine
- ✓ Architecture beyond a single server room. Simpler capacity planning
- ✓ Capacity planning only for DCI traffic flows
- ✓ Clean role separation and uniform reachability from the entire fabric are the major advantages


Border Gateway hosts:

VTEP for:

- East-West (DCI Packet Re-Originatation L2/L3)

Super-Spine hosts:

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



BGW Node Placement Option

Flexible Design option-4


BGW on Super-Spine nodes (BGW Super-Spine)

- ✓ Scale-out Multi-Clos Fabric to Interconnect the PODs using Super-Spine
- ✓ Architecture beyond a single server room. Simpler capacity planning
- ⚠ Capacity planning needs to accommodate all flows
- ⚠ Extra functional dependency (BGW + Super-Spine). **Not recommended** due to Multi-POD failure dependency

Border Gateway Super-Spine hosts:

VTEP for:


- East-West (DCI Packet Re-Originating L2/L3)
- Route Reflector (RR) (Site-Internal iBGP EVPN)
- Rendezvous Point (RP) (Site-Internal Multicast Underlay BUM)



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** ↔
- **Advertise Lo0 (Overlay Control Plane), Lo1 (BUM, External networks), and Lo100 (Multi-Site Inter-Site Transit communication)**
- 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.



VXLAN Multi-Site Architecture

DCI- BGW Back-to-Back

- Available option for connecting 2 or 3 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** $\longleftrightarrow$
 - **Advertise Lo0 (Overlay Control Plane), Lo1 (BUM, External networks), and Lo100 (Multi-Site Inter-Site Transit communication)**
 - 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** $\longleftrightarrow$
 - 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.




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

`evpn multisite border-gateway 65001`
`delay-restore time 300`



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:

$I = V \bmod N$ {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: “ $I = 48 \bmod 4 = 0 = 10.1.1.1$ ”

The DF for VLAN 49 will be: “ $I = 49 \bmod 4 = 1 = 10.1.1.2$ ”

The DF for VLAN 50 will be: “ $I = 50 \bmod 4 = 2 = 10.1.1.3$ ”


The DF for VLAN 51 will be: “ $I = 51 \bmod 4 = 3 = 10.1.1.4$ ”

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



VXLAN Multi-Site Architecture

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



Region-A


 Drop due to Designated Forwarder (DF) rule

 Physical connection

 BUM packet

VXLAN Multi-Site Architecture


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



VXLAN Multi-Site Architecture

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


- Available option for connecting 2 or 3 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** $\longleftrightarrow$
 - 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** $\longleftrightarrow$
 - 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.



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** ↔
 - Advertise **Lo0 (Overlay Control Plane)**, **Lo1 (BUM, External networks)**, and **Lo100 (Multi-Site Inter-Site Transit communication)**
 - 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.



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** $\longleftrightarrow$
 - 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** $\longleftrightarrow$
 - 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.



Special Considerations for Layer-3 extension across Multi-Site

VXLAN Native Data Plane Security




- VXLAN EVPN offers a native data plane security functionality
 - VXLAN traffic originated from a remote VTEP is only accepted when sourced from a TEP address that is an “NVE peer”
 - An NVE peer’s address is added to the local table based on the reception of MP-BGP EVPN updates carrying that specific address as next-hop
- Prevents the insertion of rogue VTEPs in a VXLAN EVPN fabric

**Why can this become an
issue for Layer-3
communication across sites?**


Inter-Site Layer 3 Traffic – Control Plane

- Inter-Site Type-2 and Type-5 EVPN updates always carry the local Multi-Site VIP as next-hop address
- Only exception are Type-5 updates for L3 networks locally connected to the BGWs



Inter-Site Layer 3 Traffic – Data Plane

Inter-Site traffic is always sourced by local BGWs from their specific PIP address
Same applies to Intra-Site traffic between local BGW and leaf nodes




The background of the slide features a dynamic, abstract pattern of flowing, wavy lines in various shades of blue, from light sky blue to deep navy blue. A prominent white rectangular box is positioned on the left side of the slide, serving as a container for the main text.

How to Prevent this?


Layer-3 Only Connectivity across Sites

Solution 1: Definition of a Loopback Interface in a Tenant VRF




Layer-3 Only Connectivity across Sites

Solution 2: Enable “advertise-pip” on the BGWs




Layer-3 Only Connectivity across Sites

Solution 3: Enable “dci-advertise-pip” on the BGWs



Installing PIP Addresses as NVE Peers

Adopting any of the 3 solutions described in the previous slide ensures that the PIP addresses of the BGWs can be installed as NVE peers in the local leaf nodes and in the remote BGW nodes



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.*

What is **Nexus Dashboard**?


Cisco Nexus Dashboard

Simple to automate, simple to consume

Powering **automation**
Unified **agile** platform



Cisco Nexus Dashboard




Automation

Accelerate provisioning
and simplify deployments



Management

In depth Management
and control for all
network deployments




Visibility

Get Centralized Visibility
and Monitoring views

Nexus Dashboard Federation

One-Manage

Multi-Cluster Nexus Dashboard Deployments





Single sign-on for clusters
(With remote authentication)



One view dashboard to see summary of fabrics, switch details etc.



VXLAN EVPN Fabric

- Unified view and automation workflows for Multi-Site deployments
- Federation of MSDs
- L2/L3 Network stretch




No. of Clusters: 10
No. of Fabrics in Multi-Cluster Group: 100
No. of Switches in Multi-Cluster: 3500

Benefits

- End-to-End connectivity and policy between loosely coupled DC sites
- Disaster Recovery, IP Mobility and Active-Active use cases
- Network policy consistency
- Change control domains separation
- Flexible geo-redundancy for ND
- Single point of Orchestration

ND-managed VXLAN Multi-Site


Use-case: Managing multiple DCs in a single instance of ND



VXLAN EVPN Multi-Site


ND-managed VXLAN Multi-Site

Use-case 1: Managing multiple DCs with multiple instances of ND



ND-managed VXLAN Multi-Site

Use-case 2: Managing single large-scale DC with 500+ switches





Automate **VXLAN Multi-Site with** **Nexus Dashboard Federation**

One-Manage: Configuration Steps

Step 1 – Create Individual ND clusters and fabrics

Step 2 – Create ND Federation across Multiple ND clusters

Step 3 – Create Multi-Cluster Fabric from ND Federation and Add Member ND/Fabrics

Step 4 – Recalculate and Deploy Multi-Site DCI Configs across BGWs/RS

Step 5 – Create, Attach, and Deploy VXLAN Overlays across Multi-Cluster Fabric

Nexus Dashboard Day-0: VXLAN EVPN Underlay

Step 1 -> Create a Fabric

Nexus Dashboard

parth-denali-1

Home

Manage

Analyze

Admin

Welcome, admin

Refresh

OverviewTopologyJourneyNewWhat's new

All Fabrics

Operational StatusConfig-Sync Status

Filter by attributes

Actions

Layout

Aggregate links

Show logical links

Create fabric

EVPN-RS

Create a new Fabric

← Topology

Create/Onboard Fabric

What is a fabric?

1 Select a category

Create new LAN fabric

2 Select a type

Greenfield (New) Fabric Option

4 Fabric summary

Brownfield (Existing) Fabric Option

Select a category

Select if you want to create a new fabric or onboard and automate configuration of an existing fabric. [Learn more](#)

Create new LAN fabric

Select this option to provision a new network comprising of Cisco NX-OS, IOS-XE, IOS-XR, and/or 3rd party devices through Nexus Dashboard.

Onboard existing LAN fabric

Select this option to preserve an existing Cisco NX-OS, IOS-XE, IOS-XR devices network's configuration and to monitor/automate the deployment of VXLAN, IP Media and Ethernet fabrics through Nexus Dashboard.

Step 1 -> Create a Fabric

Select a type

Select VXLAN

Switches in this fabric will be provisioned with the option you choose.

VXLAN

Automate a VXLAN BGP EVPN fabric for Cisco Nexus (NX-OS) and/or Catalyst (IOS-XE) switches.

Classic LAN

Automate the provisioning of a 2 or 3-tier Traditional Classical Ethernet Network.

AI/ML

Automate a Nexus (NX-OS) fabric for top performance AI/ML networks using RoCEv2.

External and inter-fabric connectivity

Monitor or manage any architecture that includes Cisco NX-OS, IOS-XE, IOS-XR and/or 3rd part devices. This includes use cases for External connectivity, Inter-fabric Connectivity Networks (such as ISNs for ACI), and Inter-Pod Networks (IPNs).

Fabric Type: Data Center VXLAN EVPN - iBGP

☒ Data Center VXLAN EVPN

Fabric for a VXLAN EVPN deployment with Nexus 9000 and/or 3000 switches.

☐ Campus VXLAN EVPN

Fabric for a VXLAN EVPN Campus deployment with Catalyst 9000 and/or Nexus 9000 switches as Border Gateways.

Settings

These are the recommended settings for configuring the parameters and capabilities of the new fabric.

Configuration mode ⓘ

Default

Advanced

Name *

DC1

Location *

Rosario, AR

BGP ASN *

65001

1-4294967295 | 1-65535[0-65535]

License tier for fabric ⓘ

☒ Essentials

☐ Advantage

☐ Premier

Enabled features

☐ Telemetry ⓘ

Cancel

Back

Next

Auto Select Name, Location, Routing Protocol ASN

Fabric Type: Data Center VXLAN EVPN - iBGP

© 2026 Cisco and/or its affiliates. All rights reserved.

BRKDCN-2988

CISCO

Step 1 -> Create a Fabric

Advanced Fields

Settings

These are the recommended settings for configuring the parameters and capabilities of the new fabric.

Configuration mode ⓘ

DefaultAdvanced

Name *

DC1

Unlock All Advanced Options

Location *

Rosario, AR

Overlay routing protocol

☒ iBGP☐ eBGP

BGP ASN *

65001


1-4294967295 | 1-65535[.0-65535]

License tier for fabric ⓘ

☒ Essentials☐ Advantage☐ Premier

Enabled features

VXLANv4 or VXLANv6



Fabric Type: Data Center VXLAN EVPN - iBGP

General ParametersReplicationvPCProtocolsSecurity

Enable IPv6 Underlay

If not enabled, IPv4 underlay is used

☐

Enable IPv6 Link-Local Address

If not enabled, Spine-Leaf interfaces will use global IPv6 addresses

☐

Fabric Interface Numbering *

p2p

Numbered(Point-to-Point) or Unnumbered

Underlay Subnet IP Mask *

30

Mask for Underlay Subnet IP Range

Underlay IP > /30 or /31

Underlay Subnet IPv6 Mask

Select an Option

Mask for Underlay Subnet IPv6 Range

Underlay Routing Protocol *

ospf

Used for Spine-Leaf Connectivity

IGP > OSPF or ISIS

Route-Reflectors *

2

Number of spines acting as Route-Reflectors

RR > 2 or 4

Anycast Gateway MAC *

2020.0000.00aa

Shared MAC address for all leafs (xxxx.xxxx.xxxx)

Distributed Anycast Gateway

© 2026 Cisco and/or its affiliates. All rights reserved.

BRKDCN-2988



Step 1 -> Create a Fabric

Advanced Fields

BUM > Multicast
or Ingress
Replication

Replication

General Parameters

vPC

Protocols

Security

Replication Mode*
multicast
Replication Mode for BUM Traffic

Multicast Group Subnet*
239.1.1.0/25
Multicast pool prefix between 8 to 30. A multicast group IP from this pool is used

IPv6 Multicast Group Subnet
IPv6 Multicast address with prefix 112 to 128

☐ Enable IPv4 Tenant Routed Multicast (TRM)
For Overlay IPv4 Multicast Support In VXLAN Fabrics

☐ Enable IPv6 Tenant Routed Multicast (TRMv6)
For Overlay IPv6 Multicast Support In VXLAN Fabrics

Default MDT IPv4 Address for TRM VRFs
Default Underlay Multicast group IPv4 address assigned for every overlay VRF

Default MDT IPv6 Address for TRM VRFs
Default Underlay Multicast group IP6 address assigned for every overlay VRF

Rendezvous-Points*
2
Number of spines acting as Rendezvous-Point (RP)

RP Mode*
asm
Multicast RP Mode. For IPv6 underlay, please use asm only

Underlay RP Loopback Id*
254
(Min:0, Max:1023)

L2VNI
Multicast
Group

RP > 2 or 4

Step 1 -> Create a Fabric

Advanced Fields

General ParametersReplicationvPCProtocolsSecurityAdvancedFreeformResources

☐ Manual Underlay IP Address Allocation

Checking this will disable Dynamic Underlay IP Address Allocations

Underlay Routing Loopback IP Range*

10.2.0.0/22

Typically Loopback0 IP Address Range

Router ID

Underlay VTEP Loopback IP Range*

10.3.0.0/22

Typically Loopback1 IP Address Range

VTEP IP

Underlay RP Loopback IP Range*

10.254.254.0/24

Anycast or Phantom RP IP Address Range

RP IP

Underlay Subnet IP Range*

10.4.0.0/16

Address range to assign Numbered and Peer Link SVI IPs

P2P Underlay IP

Layer 2 VXLAN VNI Range*30000-49000Overlay Network Identifier Range (Min:1, Max:16777214)L2VNI Label

Layer 3 VXLAN VNI Range*50000-59000Overlay VRF Identifier Range (Min:1, Max:16777214)L3VNI Label

Network VLAN Range*2300-2999Per Switch Overlay Network VLAN Range (Min:2, Max:4094)

VRF VLAN Range*2000-2299Per Switch Overlay VRF VLAN Range (Min:2, Max:4094)

VRF Lite Subnet IP Range*10.33.0.0/16Address range to assign P2P Interfabric ConnectionsVRF_LITE IP Range

VRF Lite Subnet Mask*30(Min:8, Max:31)

VRF Lite IPv6 Subnet Rangefd00::a33:0/112IPv6 Address range to assign P2P Interfabric Connections, used in L3Out

VRF Lite IPv6 Subnet Mask*126(Min:112, Max:127)Auto Allocate VRF_LITE IPs

VRF Lite Deployment*back2Back&ToExternalVRF Lite Deployment Options: 'back2Back&ToExternal' is the default. 'back2Back&ToExternal' is the default. 'back2Back&ToExternal' is the default.

☒ Auto Deploy for PeerWhether to auto generate VRF LITE sub-interface

☐ Auto Deploy Default VRFFor IPv4 underlay, whether to auto generate BGP

VRF_LITE Handoff

Edge Router Auto Deploy

☒ Auto Allocation of Unique IP on VRF Extension over VRF Lite IFCWhen enabled, IP prefix allocation is enabled on VRF Lite IFC.

☒ Per VRF Per VTEP Loopback IPv4 Auto-ProvisioningAuto provision a loopback IPv4 on a VTEP on VRF attachment. Note: VRFs are unattached.

Auto Create VTEP Loopback/EVPN Type-5 Route

Step 2 -> Add Switches

Create/Onboard Fabric



All done!

The fabric DC1, has successfully been created/onboarded to Nexus Dashboard.

You can modify any of these settings from Manage > Fabrics > Actions by clicking 'Edit Fabric Settings'

View fabric details

Add switches to fabric

Add Switches - Fabric: DC1

Switch Addition Mechanism*

☒ Discover

Seed Switch Details

Seed IP*

192.168.100.74

Switch Mgmt0 IP

Ex: "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21"

Authentication / Privacy*

MD5

Username*

parth

Switch Discovery Credentials

Max Hops*

2

Switch Hops based on CDP

Password*

.....

☐ Set as individual device write credential

© 2026 Cisco and/or its affiliates. All rights reserved.

BRKDCN-2988



Step 2 -> Add Switches

(continued) POAP option

Power-On Auto Provisioning is touchless automation of network switch bring-up

Enable Bootstrap and Local DHCP server

DHCP Scope and Default Address

General Parameters **Replication** **vPC** **Protocols** **Security** **Advanced** **Freeform** **Resources** **Manageability** **Bootstrap**

☒ **Enable Bootstrap**
Automatic IP Assignment For POAP

☒ **Enable Local DHCP Server**
Automatic IP Assignment For POAP From Local DHCP Server

DHCP Version
dhcipv4

DHCP Scope Start Address*
192.168.92.100
Start Address For Switch POAP

DHCP Scope End Address*
192.168.92.110
End Address For Switch POAP

Switch Mgmt Default Gateway*
192.168.92.1
Default Gateway For Management VRF On The Switch

Switch Mgmt IP Subnet Prefix*
24
(Min:8, Max:30)

Step 2 -> Add Switches

(continued) POAP option

Bootstrap switch in a Fabric

Switches in POAP loop

Switch Addition Mechanism*

Switch Credentials

For discovery, use*

Switches to Bootstrap

	Serial Number	Model	Version	IP Address	Hostname	Gateway	Role	Action
☐	98Z8C5LLPTL	N9K-C9300v	10.2(4)			192.168.92.1/24		Edit
☐	9REYJCVCVSI	N9K-C9500v	10.2(4)			192.168.92.1/24		Edit
☐	9UXZ1T3RXUW	N9K-C9300v	10.3(1)			192.168.92.1/24		Edit

Mgmt0 IP address

Hostname

Switch role (Access/Agg)

Edit to provide more details

Edit bootstrap switch

IP Address*

Hostname*

Image Policy

Switch Role

Data*

Cancel Save

Step 2 -> Add Switches

Discovery Results

Filter by attributes


☐	Switch Name	Serial Number	IP Address	Model	Version	Status	Progress	
☐	leaf2-dc1	9K4FDX6S4UU	192.168.100.72	N9K-C9300v	10.4(1)	In Progress		
☐	bgw1-dc1	9MJi6AJ90Y0	192.168.100.75	N9K-C9300v	10.4(1)	In Progress		
☐	leaf1-dc1	9HHS3S8BIO7	192.168.100.71	N9K-C9300v	10.4(1)	In Progress		
☐	leaf3-dc1	96NYE15D501	⚠ 192.168.100.73	N9K-C9300v	10.4(1)	In Progress		

Close

Add Switches

Switch Inventory Management

Step 3 -> Set Role



Select Role

Search Role

Spine

Leaf (current)

Border

Border Spine

Border Gateway

Border Gateway Spine

Super Spine

Border Super Spine

Border Gateway Super Spine

ToR

Switch Roles

Step 4 -> VPC Pairing

Drill-down

View details

Set role

Configuration >

Discovery >

Maintenance >

VPC pairing

ToR pairing

Delete switch

leaf1-dc1

leaf2-dc1

leaf3-dc1

Leaf VPC Pairing

vPC Pairing

Select vPC Peer for leaf1-dc1

Filter by attributes

Virtual Peerlink

Device	Recommended	Reason	Serial Number	IP Address
☒ leaf2-dc1	False	N9K-C9300v doesn't support Virtual Fabric Peering	9K4FDX6S4UU	192.168.100.72
☐ spine1-dc1	False	Switches have different roles for virtual peering	9209PCVYS2H	192.168.100.74
☐ leaf3-dc1	False	N9K-C9300v doesn't support Virtual Fabric Peering	96NYE15D501	192.168.100.73
☐ bgw1-dc1	False	Switches have different roles for virtual peering	9MJI6AJ90Y0	192.168.100.75

vPC Pairing


Select vPC Peer for leaf1-dc1

Filter by attributes

Virtual Peerlink

Device	Recommended	Reason	Serial Number	IP Address
☒ leaf2-dc1	True	Switches are connected and have same role	9K4FDX6S4UU	192.168.100.72
☐ spine1-dc1	False	Switches have different roles	9209PCVYS2H	192.168.100.74
☐ leaf3-dc1	False	Switches are not connected	96NYE15D501	192.168.100.73
☐ bgw1-dc1	False	Switches have different roles	9MJI6AJ90Y0	192.168.100.75

Step 5 -> Recalculate and Deploy



Step 5 -> Recalculate and Deploy

Deploy Configuration - DC1

1

Config Preview

2

Deploy Progress

Filter by attributes

Resync All

Switch Name	IP Address	Role	Serial Number	Fabric Status	Pending Config	Status Description	Progress	Resync Switch	
leaf3-dc1	192.168.100.73	leaf	96NYE15D501	Out-Of-Sync	575 Lines	Out-of-Sync		Resync	
leaf1-dc1	192.168.100.71	leaf	96NYE15D501	Out-Of-Sync	597 Lines	Out-of-Sync		Resync	
leaf2-dc1	192.168.100.72	leaf	9K4FDX6S4UU	Out-Of-Sync	597 Lines	Out-of-Sync		Resync	
bgw1-dc1	192.168.100.75	border gateway	9MJ16AJ90Y0	Out-Of-Sync	382 Lines	Out-of-Sync		Resync	
spine1-dc1	192.168.100.74	spine	9209PCVYS2H	Out-Of-Sync	412 Lines	Out-of-Sync		Resync	

Fabric Configuration status

Pending Configuration

Step 5 -> Recalculate and Deploy

Pending Config - DC1 - spine1-dc1

Pending Config Side-by-Side Comparison

```
feature ngoam
feature nxapi
feature ospf
feature pim
ipv6 switch-packets lla
nv overlay evpn
feature lldp
feature bgp
feature nv overlay
ip pim anycast-rp 10.254.254.1 10.2.0.5
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ngoam install acl
nxapi http port 80
nxapi https port 443
snmp-server host 172.25.74.136 traps version 2c public udp-port 2162
router bgp 65001
  router-id 10.2.0.5
  neighbor 10.2.0.1
    remote-as 65001
    update-source loopback0
  address-family l2vpn evpn
    send-community both
    route-reflector-client
  exit
  description to-leaf1-dc1
```

The diagram shows four blue boxes with arrows pointing to specific configuration lines in the pending config:

- Spine related features** points to the first five lines: `feature ngoam`, `feature nxapi`, `feature ospf`, `feature pim`, and `ipv6 switch-packets lla`.
- Spine RR/VP function** points to the line: `ip pim anycast-rp 10.254.254.1 10.2.0.5`.
- Spine RR/VP Client** points to the lines: `neighbor 10.2.0.1`, `remote-as 65001`, `update-source loopback0`, `address-family l2vpn evpn`, `send-community both`, `route-reflector-client`, and `exit`.
- Spine RR/VP function** also points to the line: `ip pim ssm range 232.0.0.0/8`.

Pending Config - DC1 - spine1-dc1

Pending Config Side-by-Side Comparison

```
interface loopback0
  ip address 10.2.0.5/32
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  description Routing loopback interface
  no shutdown
interface loopback254
  ip address 10.254.254.1/32
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  description RP loopback interface
  no shutdown
interface ethernet1/1
  no switchport
  ip address 10.4.0.5/30
  description connected-to-leaf1-dc1-Ethernet1/1
  mtu 9216
  ip router ospf UNDERLAY area 0.0.0.0
  ip ospf network point-to-point
  ip pim sparse-mode
  no shutdown
  negotiate auto
```

The diagram shows three blue boxes with arrows pointing to specific configuration lines in the pending config:

- EVPN Overlay Source Peering** points to the lines: `ip router ospf UNDERLAY area 0.0.0.0` and `ip pim sparse-mode` under `interface loopback0`.
- CDP Link & IGP configs** points to the lines: `ip router ospf UNDERLAY area 0.0.0.0`, `ip ospf network point-to-point`, `ip pim sparse-mode`, and `negotiate auto` under `interface ethernet1/1`.

Step 5 -> Recalculate and Deploy

Pending Config - DC1 - leaf1-dc1

Pending Config Side-by-Side Comparison

```
cfs eth distribute
feature dhcp
feature lacp
feature ngoam
feature nxapi
feature ospf
feature pim
feature vpc
nv overlay evpn
feature interface-vlan
feature vn-segment-vlan-based
feature lldp
feature nv overlay
feature bgp
evpn
fabric forwarding anycast-gateways
ip pim rp-address 10.254.254.1
ip pim ssm range 232.0.0.0/8
```

VTEP related features

Pending Config - DC1 - leaf1-dc1

Pending Config Side-by-Side Comparison

```
router bgp 65001
router-id 10.2.0.1
neighbor 10.2.0.5
remote-as 65001
update-source loopback0
address-family l2vpn evpn
send-community both
exit
description to-spine1-dc1
configure terminal
router ospf UNDERLAY
router-id 10.2.0.1
ip dhcp relay information option
ip dhcp relay information option vpc
```


EVPN session with Spine

Pending Config - DC1 - leaf1-dc1

Pending Config Side-by-Side Comparison

Running Config	Expected Config
	evpn
	fabric forwarding anycast-gateway-mac 2020.0000.00aa
	feature bgp
	feature dhcp
	feature interface-vlan
	feature lacp
	feature lldp
	feature ngoam
	feature nv overlay
	feature nxapi
	feature ospf
	feature pim
	feature vn-segment-vlan-based
	feature vpc
hostname leaf1-dc1	hostname leaf1-dc1
	interface Vlan3600
	description VPC-Peer-Link SVI
	ip address 10.4.0.2/30
	ip ospf network point-to-point
	ip pim sparse-mode
	ip router ospf UNDERLAY area 0.0.0.0
	mtu 9216
	no ip redirects
	no ipv6 redirects
	no shutdown
interface ethernet1/1	interface ethernet1/1
	description connected-to-spine1-dc1-Ethernet1/1
	ip address 10.4.0.6/30
	ip ospf network point-to-point

Step 5 -> Recalculate and Deploy



Nexus Dashboard Topology View

```
graph TD; bgw1-dc1 --- spine1-dc1; spine1-dc1 --- leaf1-dc1; spine1-dc1 --- leaf2-dc1; leaf1-dc1 --- leaf2-dc1;
```

✓ In-Sync

bgw1-dc1

Switch

[View details](#)

Connections

Local Port → Remote Port

Ethernet1/1 ✓ spine1-dc1 (Ethernet1/3)

IP Address192.168.100.75

Serial Number9MJ16AJ90Y0

RoleBorder Gateway

Fabric NameDC1

ModelN9K-C9300v

Version10.4(1)

ModeNormal

Connectivity Status

● OK

Uptime137 days, 14 hours, 30 minutes

✓ In-Sync

Connection

Source

Namespine1-dc1

InterfaceEthernet1/3

Destination

Namebgw1-dc1

InterfaceEthernet1/1

✓ In-Sync

Connections

[View all](#)

bgw1-dc1 → spine1-dc1

Loopback0 ✓ Loopback0


Ethernet1/1 ✓ Ethernet1/3

	Source Name	Source Interface	Destination Name	Destination Interface	Configuration Status	Protocol	State	
☐	bgw1-dc1	Loopback0	spine1-dc1	Loopback0	✓ In-Sync	BGP	Established	
☐	spine1-dc1	Ethernet1/3	bgw1-dc1	Ethernet1/1	✓ In-Sync	OSPF	FULL	


Nexus Dashboard Federation: Multi-Site DCI

Step 1 -> Create Individual ND Clusters and Fabrics

ND-1





ND-2



Step 1 -> Create Individual ND Clusters and Fabrics

ND-1



Step 2 -> Create Federation

Login to either ND-1 or ND-2 admin console to start ND Federation. In this case, we will use ND-1 as the Federation Manager/Primary.

The image displays three overlapping screenshots of the Cisco Nexus Dashboard configuration interface, illustrating the steps to create a federation.

Top Screenshot: System Settings
The 'System Settings' page is shown with tabs for 'General', 'Fabric management', 'Multi-cluster connectivity' (highlighted with a blue box), and 'Flow collection'. Below the tabs, a large blue box with a white cube icon and the text 'No Clusters Connected' is visible. A blue button labeled 'Connect Cluster' is highlighted with a blue box.

Middle Screenshot: Connect Cluster
The 'Connect Cluster' page is shown with a back arrow and the title 'Connect Cluster'. A progress indicator on the left shows three steps: '1 Select type' (highlighted with a blue box), '2 Settings', and '3 Summary'. The 'Select type' section contains two options: 'Nexus Dashboard' (highlighted with a blue box) and 'APIC'. The 'Nexus Dashboard' option is described as 'Unified platform to transform data-center network operations with simplicity, automation, and analytics.'

Bottom Screenshot: Settings
The 'Settings' page is shown with the title 'Settings' and a subtitle 'Provide the details of the remote Nexus Dashboard you are attempting to connect.' The form contains the following fields: 'Hostname/IP address' (with the value '172.25.74.149'), 'Username' (with the value 'papatel2'), 'Password' (with masked characters and a 'Show' link), and 'Login Domain' (with a dropdown menu showing 'Tacacs').

ND Federation All Cluster View

One-View

System Settings

GeneralFabric managementMulti-cluster connectivityFlow collection

Create a site to connect clusters together.

Filter by attributes

Name	Connectivity status	Version
parth-denali-1 Primary Local	Up	4.0.1i
parth-denali-2	Up	4.0.1i

Welcome, papatel2

OverviewTopology

Global One-View

SWIM
Data Center VXLAN EVPN - iBGP

ECL-G
Enhanced Classic LAN

External-Network
External and Inter-Fabric Connectivity

CE-Site-1
Enhanced Classic LAN

DC2
Data Center VXLAN EVPN - iBGP

CC
Data Center VXLAN EVPN - iBGP

Anomaly Level Critical
7 total critical anomalies, out of which 7 occurred in the last week

Advisory Level Healthy
No advisories found

Clusters2

Fabrics9

Inter-Fabric2

Switches23

ND Federation All Cluster View

One-View

Clusters

2

Nexus Dashboard Clusters

You have 2 ND clusters with a total of 9 fabrics between them. 2 of critical.

See More in Local Cluster's System Status

All Clusters

Filter by attributes

Name	Type
parth-denali-1	ND
parth-denali-2	ND

Fabrics

9

Fabrics Onboarded

You have 9 fabrics onboarded across all of your clusters. 3 Fabric(s) have an anomaly score of critical.

See More in Manage

All Fabrics

Filter by attributes

Fabric Name	Type	Cluster Name
CC	Data Center VXLAN EVPN - IBGP	parth-denali-1
CE-Site-1	Enhanced Classic LAN	parth-denali-1
DC1	Data Center VXLAN EVPN - IBGP	parth-denali-1
DC2	Data Center VXLAN EVPN - IBGP	parth-denali-2
ECL-G	Enhanced Classic LAN	parth-denali-2
EVPN-RS	External and Inter-Fabric Connectivity	parth-denali-1
External-Network	External and Inter-Fabric Connectivity	parth-denali-1
Fab1	Data Center VXLAN EVPN - IBGP	parth-denali-1
SWIM	Data Center VXLAN EVPN - IBGP	parth-denali-1

Switches

23

Switches

You have 23 switches in this clusters. All of your switches are connected to Nexus Dashboard.

See More in Manage

All Switches

Filter by attributes

Name	Fabric	Cluster	Anomaly level	Advisory level	IP Address	Model	Configura... Sync Status	Type	Role	Serial Number	Software Version	Discovery Status
S1-LEAF-1-DDD36RU38	SWIM	parth-denali-1	Major	Healthy	10.18.190.88	N9K-C93600CD-GX	InSync	Data Center VXLAN EVPN - IBGP	leaf	FDO2314199J	10.3(3)	Ok
S1-LEAF-2-DDD36RU37	SWIM	parth-denali-1	Major	Healthy	10.18.190.89	N9K-C93600CD-GX	InSync	Data Center VXLAN EVPN -	leaf	FDO23141951	10.4(1)	Ok

5 Data Center VXLAN EVPN - IBGP

Fabrics Onboarded

2 Enhanced Classic LAN

Fabrics Onboarded

2 External And Inter-Fabric Connectivity



Fabrics Onboarded

3 Critical

Highest anomaly level

Step 3 -> Create Multi-Cluster Fabric Group

VXLAN Multi-Site across Multi-Cluster ND



Step 3 -> Create Multi-Cluster Fabric Group

Fabric Group Name*

Multi-Site-Argentina

Fabric Group Type

VXLAN EVPN Multi-Site

General Parameters

DCI

Security

Resources

Configuration Backup

Layer 2 VXLAN VNI Range*

30000-49000

Overlay Network Identifier Range (Min:1, Max:16777214)

Layer 3 VXLAN VNI Range*

50000-59000

Overlay Network Identifier Range (Min:1, Max:16777214)

☐ Enable Downstream VNI

When enabled, allows different VXLAN fabric to have different VNI for the same vrf or network.

Layer 2 VXLAN VNI Global Range

Unique Range for L2VNI when Downstream VNI is enabled (Min:1, Max:16777214). This range should not have conflict with any VNI already used in member fabric

Layer 3 VXLAN VNI Global Range

Unique Range for L3VNI when Downstream VNI is enabled (Min:1, Max:16777214). This range should not have conflict with any VNI already used in member fabric

☐ Enable IPv6 Underlay for VXLAN Fabric

If not enabled, IPv4 underlay is used in child VXLAN fabric

Anycast-Gateway-MAC

2020.0000.00aa

Shared MAC address for all leafs

Multi-Site VTEP VIP Loopback Id*

100

(Min:0, Max:1023)

Border Gateway IP TAG

54321

Routing tag associated with IP address of loopback and DCI interfaces. Tag value 0-4294967295

General Parameters

DCI

Security

Resources

Configuration Backup

Multi-Site Overlay IFC Deployment Method*

Centralized_To_Route_Server

Manual, Auto Overlay EVPN Peering to Route Server, Auto Border Gateway

Multi-Site Route Server List*

9.9.9.9

Multi-Site Route-Server peer list (typically loopback IP address on Route-Server for Multi-Site EVPN peer underlay, e.g. fd00::b10:1, fd00::b10:2)

Multi-Site Route Server BGP ASN List*

65099

1-4294967295 | 1-65535[0-65535], e.g. 65000, 65001

☐ Enable 'redistribute direct' on Route Servers

For auto-created Multi-Site overlay IFCs in Route Servers. Applicable only when Multi-Site

Route Server IP T

Routing tag associated with Route Server IP for redistribute direct. This is the IP used in eBGP E

☒ BGP Neighbor Description on Multi-Site Overlay IFC

Generate BGP neighbor description for auto-created Multi-Site overlay IFCs

☒ Multi-Site Underlay IFC Auto Deployment Flag

☒ BGP Send-community on Multi-Site Underlay IFC

For auto-created Multi-Site Underlay IFCs

☐ BGP log neighbor change on Multi-Site Underlay IFC

For auto-created Multi-Site Underlay IFCs

☐ BGP BFD on Multi-Site Underlay IFC

For auto-created Multi-Site Underlay IFCs

Delay Restore time

300

Multi-Site underlay and overlay control plane convergence time (Min:30, Max:1000) in seconds

General Parameters

Security

Resources

Configuration

Multi-Site VTEP VIP Loopback IP Range*

10.10.0.0/24

Typically Loopback100 IP Address Range

DCI Subnet IP Range*

10.10.1.0/24

Address range to assign P2P DCI Links

Subnet Target Mask*

30

Target Mask for Subnet Range (Min:8, Max:31)

Stretched VNI Namespace

DSVNI Feature

Anycast Gateway

DCI Design Option

Auto IP Allocation

BGW BGP Flags

Step 4 -> Add Fabrics inside Multi-Cluster Fabric Group

Add Child fabrics (DC1, DC2, EVPN-RS) one-by-one inside the Multi-Cluster Fabric Group

Multi-Cluster Fabric Group - Multi-Site-Argentina

Refresh

View in topology

Actions

Overview

Inventory

Connectivity

Segmentation and security

Configuration Policies

History

Child Fabrics

Switches

Filter by attributes

Actions

Fabric Name	Fabric Technology	Fabric Type	ASN	Cluster
-------------	-------------------	-------------	-----	---------

Add Child Fabric

Filter by attributes

	Fabric Name	Cluster Name	Fabric Technology	Fabric Type	ASN	Fabric Health
☐	CC	parth-denali-1	VXLAN EVPN	Data Center VXLAN EVPN	65555	Critical
☐	External-Network	parth-denali-1	Custom	Custom Network	65000	Minor
☒	DC1	parth-denali-1	VXLAN EVPN	Data Center VXLAN EVPN	65001	Major
☐	EVPN-RS	parth-denali-1	Custom	Custom Network	65099	Minor
☐	Fab1	parth-denali-1	VXLAN EVPN	Data Center VXLAN EVPN	65001	Major
☐	SWIM	parth-denali-1	VXLAN EVPN	Data Center VXLAN EVPN	65001	Critical
☐	DC2	parth-denali-2	VXLAN EVPN	Data Center VXLAN EVPN	65002	Critical
☐	Test	parth-denali-2	VXLAN EVPN	VXLAN EVPN Multi-Site		Healthy

Step 5 -> Perform R&D for DCI Config

Multi-Cluster Fabric Group - Multi-Site-Argentina

Refresh

View in topology

Actions ^

Overview

Inventory

Connectivity

Segmentation and security

Configuration Policies

History

Child Fabrics

Switches

Edit Fabric Group Settings

Add Child Fabric

Recalculate and Deploy

Maintenance >

Deploy Fabric Group Configuration - Multi-Site-Argentina

1

2

Config Preview

Deploy Progress

Filter by attributes

Resync All

Cluster Name	Fabric Name	Switch Name	IP Address	Role	Serial Number	Fabric Status	Pending Config	Status Description	Progress	Resync Switch	
parth-denali-2	DC2	spine-bgw1-dc2	192.168.100.77	border gateway spine	9N19GSLKTYD	Out-Of-Sync	58 Lines	Out-of-Sync		Resync	
parth-denali-1	EVPN-RS	evpn-route-server	192.168.100.76	core router	9VSEXGYXSVN	Out-Of-Sync	48 Lines	Out-of-Sync		Resync	
parth-denali-1	DC1	bgw1-dc1	192.168.100.75	border gateway	9MJ16AJ90Y0	Out-Of-Sync	57 Lines	Out-of-Sync		Resync	

Multi-Site DCI Config Preview

Pending Config - Multi-Site-Argentina - **evpn-route-server**

Pending Config Side-by-Side Comparison

```
interface ethernet1/2
  no switchport
  ip address 10.10.1.2/30 tag 54321
  mtu 9216
  negotiate auto
  no shutdown
router bgp 65099
  address-family ipv4 unicast
    maximum-paths 64
    maximum-paths ibgp 64
    network 9.9.9.9/32
  exit
  address-family ipv6 unicast
    maximum-paths 64
    maximum-paths ibgp 64
  exit
  address-family l2vpn evpn
    retain route-target all
  exit
  template peer OVERLAY-PEERING
    update-source loopback0
    ebgp-multihop 5
    address-family l2vpn evpn
      route-map unchanged out
      send-community both
    exit
  exit
```

**DCI Underlay Routed
Port towards BGW**

**Retain RT coming from
BGW EVPN NLRI**

**EVPN NH unchanged on RS for BGW-to-
BGW VTEP termination**

```
neighbor 10.10.1.1
  remote-as 65001
  update-source Ethernet1/2
  address-family ipv4 unicast
    next-hop-self
    send-community both
  exit
  exit
neighbor 10.2.0.4
  remote-as 65001
  inherit peer OVERLAY-PEERING
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
  exit
  exit
neighbor 20.2.0.1
  remote-as 65002
  inherit peer OVERLAY-PEERING
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
```

**eBGP IPv4 Unicast (default) DCI
Underlay peering towards BGW**

**eBGP EVPN DCI Overlay peering
towards BGW**

Rewrite ASN portion of Auto RT

Multi-Site DCI Config Preview

Pending Config - Multi-Site-Argentina - bgw1-dc1

Pending Config Side-by-Side

**BGW Interface
advertisement for DCI
reachability**

```
route-map rmap-redist-direct permit 10
  match tag 54321
evpn multisite border-gateway 65001
  delay-restore time 300
router bgp 65001
  address-family ipv4 unicast
    redistribute direct route-map rmap-redist-direct
    maximum-paths 64
    maximum-paths ibgp 64
  exit
  address-family ipv6 unicast
    maximum-paths 64
    maximum-paths ibgp 64
  exit
  neighbor 10.10.1.2
    remote-as 65099
    update-source Ethernet1/2
    address-family ipv4 unicast
      next-hop-self
      send-community both
    exit
  exit
```

BGW Site ID

**eBGP IPv4 Unicast
(default) DCI Underlay
peering towards RS**

**eBGP EVPN DCI
Overlay peering
towards RS**

Pending Config

```
neighbor 9.9.9.9
  remote-as 65099
  update-source loopback0
  ebgp-multihop 5
  peer-type fabric-external
  address-family l2vpn evpn
    send-community both
    rewrite-evpn-rt-asn
  exit
configure terminal
interface nve1
  multisite border-gateway interface loopback100
interface loopback100
  ip address 10.10.0.1/32 tag 54321
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown
```

**Rewrite and Re-origination
function on BGW**

**Rewrite ASN
portion of Auto RT**

**BGW VTEP Interface
Lo1 = NVE PIP
Lo100 = NVE Multi-Site
Anycast VIP**

**BGW DCI-tracking towards
Site-External underlay**

```
interface ethernet1/2
  no switchport
  ip address 10.10.1.1/30 tag 54321
  evpn multisite dci-tracking
  mtu 9216
  negotiate auto
  no shutdown
interface ethernet1/1
  no switchport
  ip address 10.4.0.14/30
  evpn multisite fabric-tracking
  description connected-to-spine1-dc1-Ethernet1/3
  mtu 9216
  ip router ospf UNDERLAY area 0.0.0.0
  ip ospf network point-to-point
  ip pim sparse-mode
  no shutdown
  negotiate auto
```

**BGW Fabric-tracking towards
Site-Internal underlay**

Multi-Site Route-Target Problem


```
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn {65001:50000}
address-family ipv6 unicast
  route-target both auto
  route-target both auto evpn {65001:50000}
evpn
vni 30000 12
rd auto
route-target import auto {65001:30000}
route-target export auto {65001:30000}
```

```
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn {65002:50000}
address-family ipv6 unicast
  route-target both auto
  route-target both auto evpn {65002:50000}
evpn
vni 30000 12
rd auto
route-target import auto {65002:30000}
route-target export auto {65002:30000}
```



EVPN Route-Target Re-Write


The "rewrite-evpn-rt-asn" command modifies the incoming EVPN advertisements by swapping the remote AS portion in the RT with the local ASN, provided the update is coming from a neighbor that is locally configured.



Route-Server in DCI

EVPN Route-Target Re-Write by Route-Server

```
router bgp 65099
  address-family l2vpn evpn
    retain route-target all
  template peer OVERLAY-PEERING
    update-source loopback0
    ebgp-multihop 5
    address-family l2vpn evpn
      send-community
      send-community extended
      route-map unchanged out
  neighbor 20.2.0.4
    inherit peer OVERLAY-PEERING
    remote-as 65002
    address-family l2vpn evpn
      rewrite-evpn-rt-asn
```



Peer-Type Fabric-External

EVPN Split-Horizon and Route Re-Originaton


```
router bgp 65001
  neighbor 2.2.2.2
    remote-as 65002
    update-source loopback0
    ebgp-multihop 5
    peer-type fabric-external
    address-family l2vpn evpn
      send-community
      send-community extended
      rewrite-evpn-rt-asn
```

- The command is defined towards Site-External EVPN peering. It provides the capability of VXLAN packet re-origination and implements VPN split horizon mechanism.
- EVPN route coming from Site-External peer-type must not be re-advertised back into the VPN. The route is only advertised towards Site-Internal VTEPs



Failure Detection on BGWs


Use of Interface Tracking



- Because of the critical role played by the BGWs, it is critical to consider the required behavior during different failure scenarios.
- Tracking the state of the interfaces connecting the BGWs to the spines (“fabric-tracking”) and to the ISN (“dci-tracking”)
- The “dci-tracking” configuration is also required on Layer-3 interfaces locally connecting Anycast BGW nodes, needed for example, back-to-back topologies.
- Allows to define via configuration how the BGW is connected in the topology and to take proper action depending on the specific failure.
- The BGW node that gets isolated from the fabric or from the ISN needs to stop receiving traffic flows if it can not forward them to the destination.

Failure Detection on BGWs


Fabric Isolation



- The Site-Internal interfaces on BGW nodes are constantly tracked to determine their status ('evpn multisite fabric-tracking' command)

Failure Detection on BGWs


Fabric Isolation



- The Site-Internal interfaces on BGW nodes are constantly tracked to determine their status ('evpn multisite fabric-tracking' command)
- **If all the Site-Internal interfaces are detected as down on a BGW:**
 - 1) The BGW will isolate itself from Multi-Site traffic. Hence, it will shutdown Lo100 aka VIP and stop advertising to remote sites. So, the traffic from remote site will never come to this BGW as it will no longer be part of the ECMP.
 - 2) The BGW site external BGP session will be up and running but it will withdraw all EVPN routes (Type 2,3,4,5)
 - 3) The remaining BGWs withdraw all BGP EVPN Route Type 4 (Ethernet segment) routes received from the now isolated BGW because reachability is missing.
- As a result, the BGW becomes isolated from both the Site-Internal and Site-External networks.
- Seamless BGW node re-insertion using a "delay-restore" timer for the Multi-Site VIP address

Failure Detection on BGWs


DCI Isolation



- The Site-External interfaces on BGW nodes are constantly tracked to determine their status ('evpn multisite dci-tracking' command)

Failure Detection on BGWs

DCI Isolation



- The Site-External interfaces on BGW nodes are constantly tracked to determine their status ('evpn multisite dci-tracking' command)
- **If all the Site-External interfaces are detected as down on a BGW:**
 - 1) The BGW will stop advertising Multi-Site VIP aka Lo100 towards Site-Internal network. Hence, all traffic destined to remote sites will be re-routed to remaining BGWs due to ECMP.
 - 2) The isolated BGW will also withdraw EVPN Type-4 as new DF election is triggered amongst the remaining BGWs.
 - 3) The isolated BGW do not need to withdraw Type-2,3, and 5 routes as DCI interfaces and remote BGP EVPN peering are down.
- As a result, the BGW continues to operate as a Site-Internal VTEP as it's PIP remains up. Hence, if any external network is connected to this BGW, traffic from local site can still be sent to this BGW.
- Seamless BGW node reinsertion using a "delay-restore" timer for the Multi-Site VIP address

Step 6 -> Create, Attach, and Deploy VXLAN Overlays

Navigate to Multi-Cluster Fabric Group > Networks > Create

The screenshot shows the 'Create Network' form with the following fields and annotations:

- Network Name***: MyNetwork_30000. Annotation: **Auto Generated Name**.
- Layer 2 Only**: ☐. Annotation: **L2 or L3 Network**.
- VRF Name***: MyVRF_50000. Annotation: **VXLAN Overlay VRF**. A **Create VRF** button is next to the field.
- Network ID***: 30000. Annotation: **Auto Generated L2VNI Label**.
- VLAN ID**: (empty). A **Propose VLAN** button is next to the field.
- Network Template***: [Default_Network_Universal >](#)
- Network Extension Template***: [Default_Network_Extension_Universal >](#)
- General Parameters** tab is active. **IPv4 Gateway/NetMask**: 10.10.10.1/24. Annotation: **Network SVI**.

At the bottom right, there are **Close** and **Create** buttons.

Step 6 -> Create, Attach, and Deploy VXLAN Overlays

VRFs

Networks

Security Groups

Contracts

Security Associations

Protocol Definitions

Select Network

☒

Network Name

☒

MyNetwork_30000

Network ID

30000

VRF Name

MyVRF_50000

IPv4 Gateway/Prefix

10.10.10.1/24

IPv6 Gateway/Prefix

Network Status

NA

VLAN ID

2300

VLAN Name

Actions

Create

Edit

Multi-Attach

Multi-Detach

Deploy

Import

Export

Delete

Multi-Attach of Networks

1

Select Switches

2

Select Interfaces

3

Summary

Select Switches to attach all Selected Networks (1)

Total No. of Attachment : 7

Filter by attributes

☒

Cluster Name

☒

parth-denali-1

☒

parth-denali-1

☒

parth-denali-1

☒

parth-denali-2

☒

parth-denali-2

☒

parth-denali-2

☒

parth-denali-2

☒

parth-denali-2

☒

Fabric Name

☒

DC1

☒

DC1

☒

DC1

☒

DC2

☒

DC2

☒

DC2

☒

DC2

☒

Switch

☒

bgw1-dc1

☒

leaf1-dc1

☒

leaf3-dc1

☒

leaf1-dc2

☒

leaf2-dc2

☒

leaf3-dc2

☒

spine-bgw1-dc2

☒

IP Address

☒

192.168.100.75

☒

192.168.100.71

☒

192.168.100.73

☒

192.168.100.78

☒

192.168.100.79

☒

192.168.100.80

☒

192.168.100.77

☒

Serial Number

☒

9MJi6AJ90Y0

☒

9HHS3S8BIO7

☒

96NYE15D501

☒

923KAU3N5OS

☒

97X9L2HUMPW

☒

9TPRLI533YZ

☒

9N19GSLKTYD

☒

Model Number

☒

N9K-C9300v

☒

N9K-C9300v

☒

N9K-C9300v

☒

N9K-C9300v

☒

N9K-C9300v

☒

N9K-C9300v

☒

N9K-C9300v

☒

Role

☒

border gateway

☒

leaf

☒

leaf

☒

leaf

☒

leaf

☒

leaf

☒

border gateway spine

☐

Network Name

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

MyNetwork_30000

☐

Switch Name

☐

leaf1-dc1

☐

leaf1-dc2

☐

leaf3-dc1

☐

leaf3-dc2

☐

leaf2-dc2

☐

leaf2-dc2

☐

spine-bgw1-dc2

☐

Peer Switch Name

☐

leaf2-dc1

☐

☐

☐

☐

☐

☐

☐

ToR Switches

☐

☐

☐

☐

☐

☐

☐

Interfaces List

☐

leaf2-dc1(eth1/6)

☐

eth1/6

☐

eth1/4-5

☐

☐

☐

☐

☐

Action

☐

View Interfaces

☐

View Interfaces

☐

View Interfaces

☐

View Interfaces

☐

View Interfaces

☐

View Interfaces

☐

View Interfaces

© 2026 Cisco and/or its affiliates. All rights reserved.

BRKDCN-2988

CISCO

Step 6 -> Create, Attach, and Deploy VXLAN Overlays

3

Summary

✓

Select Switches

✓

Select Interfaces

Summary

Networks Selected
1

Devices Selected
7

Network Attachments
7

Device Interface Association
4

Device Interface De-Association
0

☐ Deploy Later

☒ Proceed to Full Switch Deploy(Recommended)

☐ Proceed to Individual Network Deploy

Filter by attributes

Resync All

Cluster Name	Fabric Name	Switch Name	IP Address	Role	Serial Number	Fabric Status	Pending Config	Status Description	Progress	Resync Switch	
parth-denali-2	DC2	leaf2-dc2	192.168.100.79	leaf	97X9L2HUMPW	Out-Of-S	53 Lines	Out-of-Sync		Resync	
parth-denali-2	DC2	leaf1-dc2	192.168.100.78	leaf	923KAU3N5OS	Out-Of-S	61 Lines	Out-of-Sync		Resync	
parth-denali-2	DC2	leaf3-dc2	192.168.100.80	leaf	9TPRLI533YZ	Out-Of-S	53 Lines	Out-of-Sync		Resync	
parth-denali-1	DC1	leaf3-dc1	192.168.100.73	leaf	96NYE15D501	Out-Of-S	72 Lines	Out-of-Sync		Resync	
parth-denali-1	DC1	bgw1-dc1	192.168.100.75	border gateway	9MJi6AJ90Y0	Out-Of-S	51 Lines	Out-of-Sync		Resync	
parth-denali-1	DC1	leaf2-dc1	192.168.100.72	leaf	9K4FDX6S4UU	Out-Of-S	66 Lines	Out-of-Sync		Resync	
parth-denali-1	DC1	leaf1-dc1	192.168.100.71	leaf	9HHS3S8BIO7	Out-Of-S	58 Lines	Out-of-Sync		Resync	
parth-denali-2	DC2	spine-bgw1-dc2	192.168.100.77	border gateway spine	9N19GSLKTYD	Out-Of-S	48 Lines	Out-of-Sync		Resync	

Close

Deploy All

Multi-Site Overlay Config Preview

Pending Config - Multi-Site-Argentina - bgw1-dc1

Pending Config Side-by-Side Comparison

```
vlan 2000
  vn-segment 50000
configure terminal
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
exit
```

L3VNI VRF Configs

Pending Config Side-by-Side Comparison

VXLAN Tenant

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

```
configure terminal
interface nve1
  member vni 50000 associate-vrf
  member vni 30000
  multisite ingress-replication
  mcast-group 239.1.1.1
interface loopback2
  vrf member myvrf_50000
  ip address 10.5.0.2/32 tag 12345
vlan 2300
  vn-segment 30000
configure terminal
evpn
  vni 30000 l2
  rd auto
  route-target import auto
  route-target export auto
```

NVE Tunnel Configs

L2VNI Network Configs

Multi-Site Overlay Config Preview

Pending Config - Multi-Site-Argentina - leaf1-dc2

Pending Config Side-by-Side Comparison

```
vlan 2000
  vn-segment 50000
configure terminal
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
exit
```

L3VNI VRF Configs

Pending Config Side-by-Side Comparison

```
interface Vlan2000
  vrf member myvrf_50000
  ip forward
  ipv6 address use-link-local-only
  no ip redirects
  no ipv6 redirects
  mtu 9216
  no shutdown
exit
router bgp 65002
  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
```

VXLAN Tenant

L3 SVI Configs

NVE Tunnel Configs

```
interface nve1
  member vni 50000 associate-vrf
  member vni 30000
  mcast-group 239.1.1.1
vlan 2300
  vn-segment 30000
configure terminal
interface Vlan2300
  vrf member myvrf_50000
  ip address 10.10.10.1/24 tag 12345
  fabric forwarding mode anycast-gateway
  no shutdown
exit
evpn
  vni 30000 l2
  rd auto
  route-target import auto
  route-target export auto
```


L2VNI Network Configs

```
interface ethernet1/6
  switchport
  switchport mode trunk
  mtu 9216
  spanning-tree port type edge trunk
  negotiate auto
  shutdown
  switchport trunk allowed vlan 2300
```

One-Manage Multi-Site Topology View

↑
Site-External DCI
↓


↑
Site-Internal Fabric
↓






VXLAN Multi-Site Asymmetric VNI

Network Namespace in Classic Ethernet




Overlay Namespace in VXLAN EVPN



Symmetric VNI Allocation


Overlay Namespace in VXLAN EVPN



Asymmetric VNI Allocation

Downstream VNI in VXLAN Multi-Site

Unique VNI Allocation Per-VXLAN Site



NXOS MAC RT Auto == ASN:L2VNI (4-byte BGP AS uses 23456 as 2-byte field)

NXOS IP RT Auto == ASN:L3VNI (4-byte BGP AS uses 23456 as 2-byte field)

NXOS MAC RT Auto == ASN:L2VNI (4-byte BGP AS uses 23456 as 2-byte field)

NXOS IP RT Auto == ASN:L3VNI (4-byte BGP AS uses 23456 as 2-byte field)

Downstream VNI in VXLAN Multi-Site

Nexus Dashboard Implementation

VRF overview - CISCO

Overview VRF Attachments Networks

Filter by attributes

☐	VRF name	VRF ID	VLAN ID	Model	Switch	Configuration status	Attachment	Switch Role	Fabric name	Loopback ID	Loopback IPV4	
☐	CISCO	3000	2000	9YM7600LXIJ	DC-Sydney-Leaf	DEPLOYED	Attached	leaf	DC-Sydney	2	10.5.0.1	
☐	CISCO	3000	2000	968WHIYU9C3	DC-Sydney-BGW	DEPLOYED	Attached	border gateway spine	DC-Sydney	2	10.5.0.2	

Actions Refresh

VRF CISCO in DC-Syndey corresponds to L3VNI 3000. The Auto RT will be 65800:3000 for Type-5 EVPN Routes

Asymmetric VNI

VRF overview - CISCO

Overview VRF Attachments Networks

Filter by attributes


☐	VRF name	VRF ID	VLAN ID	Model	Switch	Configuration status	Attachment	Switch Role	Fabric name	Loopback ID	Loopback IPV4 Addi	
☐	CISCO	7000	2000	9W686V5LKZ	DC-Melbourne-Leaf	DEPLOYED	Attached	leaf	DC-Melbourne	2	10.6.0.1	
☐	CISCO	7000	2000	9R1I0H1FF8V	DC-Melbourne-BGW	DEPLOYED	Attached	border gateway spine	DC-Melbourne	2	10.6.0.2	

Actions Refresh

VRF CISCO in DC-Melbourne corresponds to L3VNI 7000. The Auto RT will be 65900:7000 for Type-5 EVPN Routes

Downstream VNI in VXLAN Multi-Site

Nexus Dashboard Implementation



During VXLAN Multi-Site Fabric Group, Nexus Dashboard detects conflict across the fabrics and VXLAN Extension fails!

Downstream VNI in VXLAN Multi-Site

Nexus Dashboard Implementation

VRF	Fabric Group VNI	Sydney VNI	Melbourne VNI
CISCO	10050000	3000	7000

Edit Multi-Site-AU Settings

Name *

Multi-Site-AU

Type *

vxlan

General Parameters

DCI

Security

Resources

Configuration Backup

Layer 2 VXLAN VNI Range

Overlay Network Identifier Range (Min:1, Max:16777214)

Layer 3 VXLAN VNI Range

Overlay VRF Identifier Range (Min:1, Max:16777214)

☒ Enable Downstream VNI

When enabled, allows different VXLAN fabric to have different VNI for the same vrf or network

Layer 2 VXLAN VNI Global Range*

10030000-10049000

Unique Range for L2VNI when Downstream VNI is enabled (Min:1, Max:16777214). This range should not have conflict with any VNI already used in member fabric

Layer 3 VXLAN VNI Global Range*

10050000-10059000

Unique Range for L3VNI when Downstream VNI is enabled (Min:1, Max:16777214). This range should not have conflict with any VNI already used in member fabric

Enable the Downstream VNI Knob at the Fabric Group level before adding any child fabrics!

Downstream VNI in VXLAN Multi-Site

Nexus Dashboard Implementation

Pending Config - Multi-Site-AU - DC-Sydney-BGW

Pending Config Side-by-Side Comparison

```
configure terminal
ip extcommunity-list standard CISCO seq 10 permit rt 23456:3000
route-map ms-fabric-to-external-rmap permit 1
  match extcommunity CISCO
  set extcomm-list CISCO delete
```

EXT Community to tag and remove RT before advertising the EVPN routes to remote DC

```
vlan 2000
  name CISCO
  vn-segment 3000
vrf context cisco
  description CISCO
  vni 3000
  rd auto
```

Auto RT to match EVPN Routes with local VTEP/Leaf

```
address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn
  route-target import 23456:10050000
  route-target export 23456:10050000
  route-target import 23456:10050000 evpn
  route-target export 23456:10050000 evpn
```

Custom RT/DSVNI RT to match EVPN Routes with Remote VTEP/BGW

```
neighbor 20.2.0.2
  description multi-site overlay session to DC-Melbourne-BGW
  remote-as 65900
  update-source loopback0
  ebgp-multihop 5
  peer-type fabric-external
  address-family l2vpn evpn
  send community both
  route-map ms-fabric-to-external-rmap out
  rewrite evpn rt asn
exit
```

BGP RMAP towards Remote VTEP/BGW

Pending Config - Multi-Site-AU - DC-Melbourne-BGW

Pending Config Side-by-Side Comparison


```
configure terminal
ip extcommunity-list standard CISCO seq 10 permit rt 23456:7000
route-map ms-fabric-to-external-rmap permit 1
  match extcommunity CISCO
  set extcomm-list CISCO delete
```

```
vlan 2000
  name CISCO
  vn-segment 7000
vrf context cisco
  description CISCO
  vni 7000
  rd auto
```

```
address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn
  route-target import 23456:10050000
  route-target export 23456:10050000
  route-target import 23456:10050000 evpn
  route-target export 23456:10050000 evpn
```

```
neighbor 10.2.0.2
  description multi-site overlay session to DC-Sydney-BGW
  remote-as 65800
  update-source loopback0
  ebgp-multihop 5
  peer-type fabric-external
  address-family l2vpn evpn
  send community both
  route-map ms-fabric-to-external-rmap out
  rewrite evpn rt asn
exit
```

Asymmetric VNI Config Example – Nexus Dashboard



Verification and Validation with Nexus Dashboard

Verification through Nexus Dashboard

Keeping you away from CLI



Step 1 Verify Network and VRF attachments

Job 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 ND)

Attachment deployment status

Job execution perspective

1

2

Config Preview				
Deploy Progress				
Filter by attributes				
Switch Name	IP Address	Status	Status Description	Progress
LEAF3-DC1	192.168.101.23	SUCCESS	Deployment completed.	Executed 5 / 5
SPINE1-DC1	192.168.101.24	COMPLETED	No Commands to execute.	
SPINE2-DC1	192.168.101.25	COMPLETED	No Commands to execute.	

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	

Deployment History Tool

Commands execution perspective

Fabric Overview - External-Edge

OverviewSwitchesLinksInterfacesPoliciesEvent AnalyticsHistoryResources

Deployment History

Policy Change History

Filter by attributes

HostName	Entity Name	Entity Type	Source	Commands	Status	Serial Number	Status
Edge-router	Vlan100	INTERFACE	UNDERLAY	Detailed History	● SUCCESS	91UP402KI8R	Successfully deployedadminMay 02 2023, 12:20:20
Edge-router	Vlan100	INTERFACE	UNDERLAY	Detailed History	● FAILED	91UP402KI8R	
Edge-router	SWITCH	SWITCH	UNDERLAY	Detailed History	● SUCCESS	91UP402KI8R	

Command Execution Details for Edge-router(91UP402KI8R)

Config	Status	CLI Response
vrf member eng	● SUCCESS	Warning: Deleted all L3 config on interface Vlan100 VRF eng does not exist. Create vrf to make interface Vlan100 operational Failed to configure hsrp version 2 on Vlan100 due to invalid command. Failed to execute this command, since command#7 (hsrp version 2) failed. Failed to execute this command, since command#7 (hsrp version 2) failed.
ip address 4.4.4.1/24	● SUCCESS	
no ip redirects	● SUCCESS	
no ipv6 redirects	● SUCCESS	
hsrp version 2	● FAILED	
hsrp 44	● FAILED	
ip 4.4.4.2	● FAILED	

CLI response messages for easier troubleshooting

Command Execution Details for Edge-router(91UP402KI8R)

Config	Status	CLI Response
router bgp 65100	● SUCCESS	
address-family ipv4 unicast	● SUCCESS	
network 100.1.1.0/24	● SUCCESS	

Show Commands Tool

Switch Config perspective

Switches

Fabric Name == DC1

Switch	IP Address	Role	Serial
BGW1-DC1	192.168.101.26	Gateway	9S92
LEAF1-DC1	192.168.101.21	Leaf	97EV
LEAF2-DC1	192.168.101.22	Leaf	988K
LEAF3-DC1	192.168.101.23	Leaf	9V7C

50 Rows

Actions

Change Mode
Provision RMA
Change Serial Number
Copy Run Start
Reload
Restore Switch
Show Commands
Exec Commands
Delete Switch(es)

Add Switches
Preview
Deploy
Discovery
Set Role
vPC Pairing
ToR/Access Pairing
vPC Overview
More

Commands*

show
current_running_config
show
show_bgp_evpn_neighbors
show_bgp_l2vpn_evpn_summary
show_bgp_sessions
show_capture_elam
show_epld_boot_info
show_mac_and_arp
show_nve_state

ND pre-built commands or user commands

```
1 #show ip ospf nei
2
3 OSPF Process ID UNDERLAY VRF default
4 Total number of neighbors: 2
5 Neighbor ID    Pri State          Up Time  Address      Interface
6 10.2.0.6       1 FULL/ -       2d22h    10.4.0.29    Eth1/2
7 10.2.0.1       1 FULL/ -       2d22h    10.4.0.21    Eth1/3
8
9 #show bgp l2vpn evpn su
10 BGP summary information for VRF default, address family L2VPN EVPN
11 BGP router identifier 10.2.0.3, local AS number 65001
12 BGP table version is 29, L2VPN EVPN config peers 2, capable peers 2
13 9 network entries and 16 paths using 2436 bytes of memory
14 BGP attribute entries [16/2752], BGP AS path entries [1/10]
15 BGP community entries [0/0], BGP clusterlist entries [6/24]
16
17 Neighbor      V  AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
18 10.2.0.1      4 65001  4252   4240    29    0   0   2d22h  5
19 10.2.0.6      4 65001  4252   4239    29    0   0   2d22h  5
20
21 #show int nve1
22 nve1 is up
23 admin state is up, Hardware: NVE
24 MTU 9216 bytes
25 Encapsulation VXLAN
```

Appendix

VXLAN BGP EVPN


<https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/white-paper-listing.html>

Nexus Dashboard

<https://www.cisco.com/c/en/us/products/cloud-systems-management/prime-data-center-network-manager/white-paper-listing.html>

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.
- ND provides simplified mechanism to extend and provide end-to-end network and policy consistency across regions, all with a single point of orchestration.
- ND provides a single plane of glass solution to automate and manage Nexus and Non-Nexus devices.

Complete your session surveys



Complete your surveys in the Cisco Events App.



Complete a minimum of 4 session surveys and the overall event survey to receive a unique Cisco Live t-shirt.
(from 11:30 on Thursday, while supplies last)

Continue your education



Visit the Cisco Showcase for related demos



Book your one-on-one Meet the Engineer meeting

Visit the Technical Solutions Clinics to discuss your technical questions



Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs



Visit the On-Demand Library for more sessions at CiscoLive.com/On-Demand

Thank you

CISCO Live !

