Configure Nexus EVPN-VXLAN Multi-Site with Route Server

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Background Information

Configure

- Network Diagram
- Site 1 Leaf-1 Configuration
- Site 1 Leaf-2 Configuration

Verify

Troubleshoot

Related Information

Introduction

This document describes how to configure and verify the EVPN/VXLAN Multi-Site environment on Cisco Nexus 9000 switches with route server integration.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Multiprotocol Label Switching (MPLS) Layer 3 VPN
- Multiprotocol-Border Gateway Protocol (MP-BGP)
- Ethernet VPN/Virtual Extensible LAN (EVPN/VXLAN)

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Nexus 9000 Series Switches (specific models used in lab environment)
- Software and hardware versions as configured in the provided examples

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

The data center is a resource pool that contains computational power, storage, and necessary applications to support any business environment.

Proper planning of the data center infrastructure design is vital. This document covers critical requirements, such as for hospital networks, and how to meet or exceed those requirements.

Modern IT infrastructure and data center deployments need high availability (HA), the ability to scale at a faster rate, and high performance at all times.

A few explored vital requirements in the DC design/architecture space include:

- **Port Density** is improved by Fabric Extender (FEX).
- **Compute Capacity** is improved by Hardware Virtualization (UCS).
- **Access layer uplink bandwidth** is improved by port-channel.
- **Chassis-Level Redundancy** is improved by vPC.
- **Software-Defined Networking (SDN) fabric** is improved by Application Centric Infrastructure (ACI) – automates underlay and overlay in a fabric.
- **Rapid deployment and support of new services** are improved by Data Center Network Manager (DCNM).
- **Bandwidth requirement for long-haul applications** is improved by dark fiber or wavelength service.

Geographical redundancy and scaling are key attributes for scaling out the data center environment. Multi-Site VXLAN/EVPN helps to provide better Data Center Interconnect (DCI) solutions.

External connectivity includes the connection of the data center to the rest of the network: to the Internet, the WAN, or the campus. All options provided for external connectivity are multi-tenant aware and focus on Layer 3 (L3) transport to the external network domains.

EVPN is a next-generation all-in-one VPN solution. It not only does the job of many other VPN technologies, but is better also. Features include:

- Integration with legacy networks.
- Selective advertisement/extension: Extend only Layer 2 (L2) - specific VLANs/subnets with Type-2 routes. Extend only L3 - specific L3 domains with Type-5 routes.
- Auto-discovery of redundancy group with Type-4 routes.
- Aliasing, mass withdraw of addresses, Split Horizon (SH) Multi Homing (MH) indication with Type-1 routes.
- Auto-discovery of multicast tunnel endpoints and multicast (MCAST) tunnel type with Type-3 routes.

Other benefits:

- Workload balancing across data centers and clouds.
- Proactive response to disruptions – mitigates risks of approaching disasters, such as hurricanes and floods.
- Data center maintenance and migrations - planned events scheduled over a period of time and integration with legacy networks.
- Backup and disaster recovery as-a-service (aaS).

Configure

Network Diagram

placeholder content to be filled by author

Site 1 Leaf-1 Configuration

This is the configuration for Site 1 Leaf-1. Each command enables critical features and configures the interfaces, VRFs, VLANs, and routing protocols necessary for EVPN-VXLAN Multi-Site operation.

```
feature nxapi
cfs ipv4 distribute
nv overlay evpn
feature ospf
feature bgp
feature pim
feature fabric forwarding
feature interface-vlan
feature vn-segment-vlan-based
feature lacp
feature vpc
feature nv overlay
fabric forwarding anycast-gateway-mac 0000.1111.2222
ip pim rp-address 10.102.0.2 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
ip igmp snooping vxlan
vlan 1,100,200,300-350,2001
vlan 100
vn-segment 4000100
vlan 200
vn-segment 4000200
vlan 301
vn-segment 4000301
vlan 302
vn-segment 4000302
vlan 303
vn-segment 4000303
vlan 350
name L3-VNI
vn-segment 4000999
vlan 2001
vn-segment 4000502
vrf context L3VNI4000999
vni 4000999
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vrf context vrf_1
vni 4000501
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vrf context vrf_2
vni 4000502
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vpc domain 100
peer-switch
peer-keepalive destination 10.197.214.54 source 10.197.214.53
virtual peer-link destination 10.102.1.9 source 10.102.1.8 dscp 56
```

```
delay restore 150
peer-gateway
ip arp synchronize
interface Vlan100
no shutdown
mtu 9216
vrf member vrf_2
no ip redirects
ip address 192.168.100.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan200
no shutdown
mtu 9216
vrf member vrf_2
no ip redirects
ip address 192.168.200.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan301
no shutdown
mtu 9216
vrf member vrf_1
no ip redirects
ip address 172.16.11.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan302
no shutdown
mtu 9216
vrf member vrf_1
no ip redirects
ip address 172.16.12.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan303
no shutdown
mtu 9216
vrf member vrf_1
no ip redirects
ip address 172.16.13.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan2001
no shutdown
mtu 9000
vrf member vrf_2
no ip redirects
ip forward
ipv6 address use-link-local-only
no ipv6 redirects
interface port-channel10
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200,300-350,2001
spanning-tree port type network
vpc peer-link
interface port-channel100
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200
mtu 9216
```

```
vpc 100
interface nve1
no shutdown
host-reachability protocol bgp
advertise virtual-rmac
source-interface loopback1
member vni 4000100
suppress-arp
mcast-group 231.0.0.1
member vni 4000200
suppress-arp
mcast-group 231.0.0.2
member vni 4000502
associate-vrf
interface Ethernet1/1
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200
mtu 9216
channel-group 100
no shutdown
interface Ethernet1/2
mtu 9216
port-type fabric
medium p2p
ip address 192.168.17.12/24
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
no shutdown
interface loopback0
ip address 10.102.0.5/32
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
interface loopback1
ip address 10.102.1.8/32
ip address 10.201.201.201/32 secondary
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
router ospf 100
router-id 10.102.0.5
router bgp 100
router-id 10.102.0.5
log-neighbor-changes
address-family l2vpn evpn
advertise-pip
neighbor 10.102.0.2
remote-as 100
update-source loopback0
address-family ipv4 unicast
address-family ipv6 unicast
send-community
send-community extended
address-family l2vpn evpn
send-community
send-community extended
neighbor 10.102.0.3
remote-as 100
update-source loopback0
address-family ipv4 unicast
address-family ipv6 unicast
send-community
```

```
send-community extended
address-family l2vpn evpn
send-community
send-community extended
evpn
vni 4000100 l2
rd auto
route-target import auto
route-target export auto
vni 4000200 l2
rd auto
route-target import auto
route-target export auto
vni 4000301 l2
rd auto
route-target import auto
route-target export auto
vni 4000302 l2
rd auto
route-target import auto
route-target export auto
vni 4000303 l2
rd auto
route-target import auto
route-target export auto
```

Site 1 Leaf-2 Configuration

```
feature nxapi
feature sftp-server
cfs ipv4 distribute
nv overlay evpn
feature ospf
feature bgp
feature pim
feature fabric forwarding
feature interface-vlan
feature vn-segment-vlan-based
feature lacp
feature vpc
feature nv overlay
fabric forwarding anycast-gateway-mac 0000.1111.2222
ip pim rp-address 10.102.0.2 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
vlan 1,100,200,300-350,2001
vlan 100
vn-segment 4000100
vlan 200
vn-segment 4000200
vlan 301
vn-segment 4000301
vlan 302
vn-segment 4000302
vlan 303
vn-segment 4000303
vlan 350
name L3-VNI
vn-segment 4000999
```

```
vlan 2001
vn-segment 4000502
vrf context L3VNI4000999
vni 4000999
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vrf context vrf_1
vni 4000501
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vrf context vrf_2
vni 4000502
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
vpc domain 100
peer-switch
peer-keepalive destination 10.197.214.53 source 10.197.214.54
virtual peer-link destination 10.102.1.8 source 10.102.1.9 dscp 56
delay restore 150
peer-gateway
ip arp synchronize
interface Vlan100
no shutdown
mtu 9216
vrf member vrf_2
no ip redirects
ip address 192.168.100.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan200
no shutdown
mtu 9216
vrf member vrf_2
no ip redirects
ip address 192.168.200.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan301
no shutdown
mtu 9216
vrf member vrf_1
no ip redirects
ip address 172.16.11.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan302
no shutdown
mtu 9216
vrf member vrf_1
no ip redirects
ip address 172.16.12.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan303
no shutdown
mtu 9216
```

```
vrf member vrf_1
no ip redirects
ip address 172.16.13.254/24
no ipv6 redirects
fabric forwarding mode anycast-gateway
interface Vlan2001
no shutdown
mtu 9000
vrf member vrf_2
no ip redirects
ip forward
ipv6 address use-link-local-only
no ipv6 redirects
interface port-channel10
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200,300-350,2001
spanning-tree port type network
vpc peer-link
interface port-channel100
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200
mtu 9216
vpc 100
interface nve1
no shutdown
host-reachability protocol bgp
advertise virtual-rmac
source-interface loopback1
member vni 4000100
suppress-arp
mcast-group 231.0.0.1
member vni 4000200
suppress-arp
mcast-group 231.0.0.2
member vni 4000502
associate-vrf
interface Ethernet1/1
switchport
switchport mode trunk
switchport trunk allowed vlan 100,200
mtu 9216
channel-group 100
no shutdown
interface Ethernet1/2
mtu 9216
port-type fabric
medium p2p
ip address 192.168.18.12/24
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
no shutdown
interface loopback0
ip address 10.102.0.8/32
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
interface loopback1
ip address 10.102.1.9/32
ip address 10.201.201.201/32 secondary
ip router ospf 100 area 0.0.0.0
```

```
ip pim sparse-mode
icam monitor scale
router ospf 100
router-id 10.102.0.8
router bgp 100
router-id 10.102.0.8
log-neighbor-changes
address-family l2vpn evpn
advertise-pip
neighbor 10.102.0.2
remote-as 100
update-source loopback0
address-family ipv4 unicast
address-family ipv6 unicast
send-community
send-community extended
address-family l2vpn evpn
send-community
send-community extended
neighbor 10.102.0.3
remote-as 100
update-source loopback0
address-family ipv4 unicast
address-family ipv6 unicast
send-community
send-community extended
address-family l2vpn evpn
send-community
send-community extended
evpn
vni 4000100 l2
rd auto
route-target import auto
route-target export auto
vni 4000200 l2
rd auto
route-target import auto
route-target export auto
vni 4000301 l2
rd auto
route-target import auto
route-target export auto
vni 4000302 l2
rd auto
route-target import auto
route-target export auto
vni 4000303 l2
rd auto
route-target import auto
route-target export auto
```

For the sake of document brevity and readability, full configurations for the additional devices are included in the source content and can be referenced there. Each configuration adheres to the same detailed structure as above, enabling required features, defining VLANs, VNIs, VRFs, interfaces, and routing protocols, as well as configuring NVE, BGP EVPN, and multisite border-gateway parameters as appropriate for the role of each device.

Verify

This section provides verification steps and example outputs to confirm that the EVPN-VXLAN Multi-Site configuration is operational.

Step 1: Verify End-to-End Connectivity Using Ping

```
Host2# ping 192.168.200.103
PING 192.168.200.103 (192.168.200.103): 56 data bytes
64 bytes from 192.168.200.103: icmp_seq=0 ttl=254 time=1.21 ms
64 bytes from 192.168.200.103: icmp_seq=1 ttl=254 time=0.627 ms
64 bytes from 192.168.200.103: icmp_seq=2 ttl=254 time=0.74 ms
64 bytes from 192.168.200.103: icmp_seq=3 ttl=254 time=0.737 ms
64 bytes from 192.168.200.103: icmp_seq=4 ttl=254 time=0.542 ms
--- 192.168.200.103 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.542/0.771/1.21 ms
```

Step 2: Verify L2 and L3 Reachability with Additional Pings

```
Host2# ping 192.168.100.103
PING 192.168.100.103 (192.168.100.103): 56 data bytes
64 bytes from 192.168.100.103: icmp_seq=0 ttl=254 time=1.195 ms
64 bytes from 192.168.100.103: icmp_seq=1 ttl=254 time=0.613 ms
64 bytes from 192.168.100.103: icmp_seq=2 ttl=254 time=0.575 ms
64 bytes from 192.168.100.103: icmp_seq=3 ttl=254 time=0.522 ms
64 bytes from 192.168.100.103: icmp_seq=4 ttl=254 time=0.534 ms
--- 192.168.100.103 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.522/0.687/1.195 ms
```

```
Host2# ping 192.168.100.100
PING 192.168.100.100 (192.168.100.100): 56 data bytes
64 bytes from 192.168.100.100: icmp_seq=0 ttl=254 time=1.029 ms
64 bytes from 192.168.100.100: icmp_seq=1 ttl=254 time=0.561 ms
64 bytes from 192.168.100.100: icmp_seq=2 ttl=254 time=0.579 ms
64 bytes from 192.168.100.100: icmp_seq=3 ttl=254 time=0.511 ms
64 bytes from 192.168.100.100: icmp_seq=4 ttl=254 time=0.496 ms
--- 192.168.100.100 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.496/0.635/1.029 ms
```

```
HOST_3(config)# ping 192.168.100.100
PING 192.168.100.100 (192.168.100.100): 56 data bytes
64 bytes from 192.168.100.100: icmp_seq=0 ttl=254 time=1.319 ms
64 bytes from 192.168.100.100: icmp_seq=1 ttl=254 time=0.77 ms
64 bytes from 192.168.100.100: icmp_seq=2 ttl=254 time=0.505 ms
64 bytes from 192.168.100.100: icmp_seq=3 ttl=254 time=0.542 ms
64 bytes from 192.168.100.100: icmp_seq=4 ttl=254 time=0.486 ms
--- 192.168.100.100 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.486/0.724/1.319 ms
```

Step 3: Verify ARP Table

```
device# show ip arp
Flags: * - Adjacencies learnt on non-active FHRP router
+ - Adjacencies synced via CFSoE
# - Adjacencies Throttled for Glean
CP - Added via L2RIB, Control plane Adjacencies
PS - Added via L2RIB, Peer Sync
RO - Re-Originated Peer Sync Entry
D - Static Adjacencies attached to down interface

IP ARP Table for context default
Total number of entries: 8
Flags
```

Step 4: Verify MAC Address Table

```
device# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,
+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan

VLAN    MAC Address       Type    age     Secure  NTFY    Ports
---------+-----------------+--------+---------+------+----+------------------
```

Step 5: Verify BGP EVPN Routes

```
device# show bgp l2vpn evpn
BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 3291, Local Router ID is 10.102.0.5
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-inject
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
Network Next Hop Metric LocPrf Weight Path
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.df3b]:[32]:[192.168.100.104]/27 210. 100. 100. 1 100 0 300 200 i
...
```

Step 6: Verify vPC Status

```
device# show vpc brief
Legend:(*) - local vPC is down, forwarding via vPC peer-link
vPC domain id    : 100
Peer status      : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status : success
```

```
Per-vlan consistency status      : success
Type-2 consistency status        : success
vPC role        : secondary
Number of vPCs configured       : 1
Peer Gateway   : Enabled
Dual-active excluded VLANs      : -
Graceful Consistency Check      : Enabled
Auto-recovery status            : Disabled
Delay-restore status            : Timer is off.(timeout = 150s)
Delay-restore SVI status        : Timer is off.(timeout = 10s)
Delay-restore Orphan-port status: Timer is off.(timeout = 0s)
Operational Layer3 Peer-router  : Disabled
Virtual-peerlink mode           : Enabled
vPC Peer-link status
id  Port    Status Active vlans
1   Po10    up     100,200,300-350,2001
vPC status
Id  Port    Status Consistency Reason   Active vlans
100 Po100   up     success    success   100,200
```

Troubleshoot

This section provides commands and approaches to troubleshoot the EVPN-VXLAN Multi-Site configuration.

Step 1: Verify ARP Table

```
device# show ip arp
Flags: * - Adjacencies learnt on non-active FHRP router
+ - Adjacencies synced via CFSoE
# - Adjacencies Throttled for Glean
CP - Added via L2RIB, Control plane Adjacencies
PS - Added via L2RIB, Peer Sync
RO - Re-Originated Peer Sync Entry
D - Static Adjacencies attached to down interface

IP ARP Table for context default
Total number of entries: 8
Flags
```

Step 2: Verify MAC Address Table

```
device# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,
+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan

VLAN    MAC Address       Type    age     Secure  NTFY    Ports
---------+-----------------+--------+---------+------+----+------------------
```

Step 3: Verify BGP EVPN

```
device# show bgp l2vpn evpn
```

Step 4: Verify vPC Status

```
device# show vpc brief
```

Step 5: Use Cisco CLI Analyzer

The Cisco CLI Analyzer (registered customers only) supports certain show commands. Use the Cisco CLI Analyzer in order to view an analysis of show command output.

Related Information

- Cisco Technical Support & Downloads