

VXLAN BGP EVPN Challenge Lab

[« VXLAN BGP EVPN Layer 2 Fabric | ACI Fabric Initialization »](#)

Last updated: April 25, 2017

VXLAN BGP EVPN w/ Layer-3 and Inter-VRF Routing

 Note:

Before proceeding, ensure that the Nexus 5600 Leaf switches are operating in Store-and-Forward mode. This is required for VXLAN encapsulation support on the Nexus 5600s. This command is not required on the 7K/9K platforms. This command requires a reload before taking effect.

```
hardware ethernet store-and-fwd-switching
Copy run start
reload
```

*Disable all East/West Connectivity between the Spines

*Disable all East/West Connectivity between the Leafs

- Enable the necessary features
 - The following functionality will be required (Spine switches only require a subset of these features, do not enable more than is necessary):
 - VLAN Interfaces
 - OSPF
 - BGP
 - PIM
 - BFD
 - EVPN Address-family
 - VXLAN encapsulation
 - VXLAN troubleshooting tools
 - Distributed Anycast Gateway
 - NXAPI
 - Mapping of VLANs to VNIDs
- Build Underlay Topology
 - Provision Loopback0 as the Overlay Loopback
 - Provision Loopback 1 as the Underlay Loopback
 - Do not assign any IP addresses on the links between the Spine and Leaf switches
 - Leverage OSPF as the Underlay IGP; All links should be in Area 0
 - All OSPF Router-IDs should be the Underlay Loopback address
 - No Network LSAs should exist within the LSDB
 - Implement PIM in the Underlay to enable the Multicast distribution of BUM traffic within the Overlay
 - The creation of an additional loopback on the Spines is acceptable for this task
 - Ensure redundancy is in place for the RP
 - There should not be any (S,G) state on any switches in the fabric
 - Implement BFD for fast failure detection in the underlay for both OSPF and PIM
 - Once the underlay is completed, all devices should be able to ping both Underlay and Overlay Loopback interfaces
- Create VLANs and VNIDs
 - On the appropriate switches, provision the following (VLAN and VNID values are based on Rack Number, the numbers below are based on Rack-3):
 - VLANs 103, 203, 2031 and 2032
 - VLAN 103 should map to VNID 10103
 - VLAN 203 should map to VNID 10203
 - VLAN 2031 should map to VNID 12031
 - VLAN 2032 should map to VNID 12032
 - VRFs RED and BLUE
 - Each VRF should be enabled for the ipv4 address-family, as well as EVPN
 - Route-Targets and Route-Distinguishers should be automatically generated
 - VRF RED L3VNI is 12031
 - VRF BLUE L3VNI is 12032
 - Server-5 should be connected to N5K5 in VLAN 103 and configured as follows:
 - IP address 10.103.103.100/24
 - MAC address: aaaa.aaaa.aa01
 - Gateway 10.103.103.1
 - Server-6 should be connected to N5K6 in VLAN 203 and configured as follows:
 - IP address: 10.203.203.100/24
 - MAC address: cccc.cccc.cc01
 - Gateway 10.203.203.1
 - Interfaces where servers are connected should immediately transition to a STP forwarding state
 - VLAN interfaces should be created on the appropriate switches for both VLAN 103 and 203
 - Server-5 and Server-6 are VMs, and frequently move between hypervisor hosts. These servers should never be required to Re-ARP for their Gateway's MAC address or need to modify their Gateway IP address, no matter what Leaf their respective hypervisor is connected to.
 - Do not use any FHRPs to accomplish this
- Establish Overlay BGP Peering and Configuration
 - All devices within the fabric will operate within BGP Autonomous System 65000
 - All BGP router-IDs should be the Overlay Loopback address
 - Establish BGP adjacencies in the required address-families i. Leaf switches must not peer with each other ii. All BGP adjacencies must be established using the OVERLAY Loopback iii. Bonus: Implement BGP on the Spines using no more than 10 lines of BGP configuration
 - Ensure communities are included in BGP Updates for the appropriate address families
 - Any IPv4 Unicast Prefixes within the RED VRF should be imported into L2VPN updates, and advertised via the L2VPN address-family
 - Any IPv4 Unicast Prefixes within the BLUE VRF should be imported into L2VPN updates, and advertised via the L2VPN address-family

- Do not establish any BGP adjacencies in the IPv4 Unicast address-family
- Implement a protection mechanism to ensure that only the desired prefixes are placed into the local IPv4 Unicast BGP table for VRF RED and VRF BLUE
- Configure EVPN L2VNI's on the required switches i. Route-Distinguishers and Route-Targets should be automatically generated
- Creation of VXLAN Overlay Interface
 - On each VTEP switch, configure the interface required for VXLAN tunnel sourcing and termination
 - The Overlay Loopback IP should be used for addressing of this interface
 - BGP must be used as the control-plane for end-host advertisement and reachability
 - Add the appropriate L2VNIs to this interface
 - VNID 103 should use 239.103.103.103 for BUM traffic forwarding
 - VNID 203 should use 239.203.203.203 for BUM traffic forwarding
 - Map the L3VNIs to their respective VRF
 - ARP requests for hosts that are already known to the EVPN fabric should not be flooded across the fabric
- Inter-VRF routing
 - On the Border-Leaf (Higher numbered 5K), ensure the backbone router can be reached in VRF RED and VRF BLUE
 - Backbone router is attached to Border-Leaf on E1/15
 - Connectivity to the Backbone router is accomplished through the use of subinterfaces
 - IP addressing for E1/15 subinterfaces is based up Rack Number and Border-Leaf switch ID
 - For example, Rack-3:
 - Border-Leaf is N5K6
 - Int e1/15.55 is in VRF RED, IP address of 10.0.55.56/24
 - Int e1/15.56 is in VRF BLUE, IP address of 10.0.56.56/24
 - Backbone router IP addresses are 10.0.55.1/24 and 10.0.56.1/24 respectively
 - On the Border-Leaf, create a Default Route for VRF RED and one for VRF BLUE; the next-hop of this Default Route should be the Backbone router
 - Advertise this newly created Default Route into each VRF via BGP
- Verify
 - Server-5 should now be able to ping Server-6

Configuration [Click to collapse](#)

```

N5K5:
install feature-set fabric
feature-set fabric
hostname N5K5
!
feature fabric forwarding
nv overlay evpn
feature ospf
feature bgp
feature pim
feature interface-vlan
feature bfd
feature nv overlay
feature nxapi
feature vn-segment-vlan-based
!
feature ngoam
!
hardware ethernet store-and-fwd-switching
!
fabric forwarding anycast-gateway-mac AAAA.BBBB.CCCC
ip pim rp-address 7.7.7.7 group-list 224.0.0.0/4 bidir
ip pim ssm range 232.0.0.0/8
ip pim bfd
!
vlan 103
  vn-segment 10103
vlan 203
  vn-segment 10203
vlan 2031
  vn-segment 12031
vlan 2032
  vn-segment 12032
!
route-map RM-BLUE-ROUTES permit 10
  match tag 2032
route-map RM-RED-ROUTES permit 10
  match tag 2031
!
vrf context BLUE
  vni 12032
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
vrf context RED
  vni 12031
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
!
interface Vlan103
  no shutdown
  vrf member RED
  ip address 10.103.103.1/24 tag 2031
  fabric forwarding mode anycast-gateway
!
interface Vlan203
  no shutdown
  vrf member BLUE
  ip address 10.203.203.1/24 tag 2032
  fabric forwarding mode anycast-gateway

```

```

!
interface Vlan2031
  description L3 VNI for VRF RED
  no shutdown
  vrf member RED
  ip forward
!
interface Vlan2032
  description L3 VNI for VRF BLUE
  no shutdown
  vrf member BLUE
  ip forward
!
interface nve1
  no shutdown
  source-interface loopback0
  host-reachability protocol bgp
  member vni 10103
    suppress-arp
    mcast-group 239.103.103.103
  member vni 10203
    suppress-arp
    mcast-group 239.203.203.203
  member vni 12031 associate-vrf
  member vni 12032 associate-vrf
!
interface Ethernet1/1
  description SRV5
  switchport access vlan 103
  spanning-tree port type edge
!
interface Ethernet1/5
  no switchport
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  no ip redirects
  ip unnumbered loopback1
  ip ospf network point-to-point
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
!
interface Ethernet1/6
  no switchport
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  no ip redirects
  ip unnumbered loopback1
  ip ospf network point-to-point
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
!
interface loopback0
  description OVERLAY LOOPBACK
  ip address 2.2.2.55/32
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
!
interface loopback1
  description UNDERLAY LOOPBACK
  ip address 1.1.1.55/32
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
!
router ospf UNDERLAY
  bfd
  router-id 1.1.1.55
!
router bgp 65000
  router-id 2.2.2.55
  address-family l2vpn evpn
  template peer SPINE_PEERS
    update-source loopback0
  address-family l2vpn evpn
    send-community both
  neighbor 2.2.2.75 remote-as 65000
    inherit peer SPINE_PEERS
    description 7K5
  neighbor 2.2.2.76 remote-as 65000
    inherit peer SPINE_PEERS
    description 7K6
  vrf BLUE
    address-family ipv4 unicast
      advertise l2vpn evpn
      redistribute direct route-map RM-BLUE-ROUTES
  vrf RED
    address-family ipv4 unicast
      advertise l2vpn evpn
      redistribute direct route-map RM-RED-ROUTES
!
!
evpn
  vni 10103 12
    rd auto
    route-target import auto
    route-target export auto
  vni 10203 12
    rd auto
    route-target import auto
    route-target export auto
!
ngoam install acl

```

```

install feature-set fabric
feature-set fabric
hostname N5K6
!
feature fabric forwarding
nv overlay evpn
feature ospf
feature bgp
feature pim
feature interface-vlan
feature bfd
feature nv overlay
feature nxapi
feature vn-segment-vlan-based
!
feature ngoam
!
hardware ethernet store-and-fwd-switching
!
fabric forwarding anycast-gateway-mac AAAA.BBBB.CCCC
ip pim rp-address 7.7.7.7 group-list 224.0.0.0/4 bidir
ip pim ssm range 232.0.0.0/8
ip pim bfd
!
vlan 103
    vn-segment 10103
vlan 203
    vn-segment 10203
vlan 2031
    vn-segment 12031
vlan 2032
    vn-segment 12032
!
route-map RM-BLUE-ROUTES permit 10
    match tag 2032
route-map RM-RED-ROUTES permit 10
    match tag 2031
!
vrf context BLUE
    vni 12032
    ip route 0.0.0.0/0 10.0.56.1 tag 2032
    rd auto
    address-family ipv4 unicast
        route-target both auto
        route-target both auto evpn
vrf context RED
    vni 12031
    ip route 0.0.0.0/0 10.0.55.1 tag 2031
    rd auto
    address-family ipv4 unicast
        route-target both auto
        route-target both auto evpn
!
interface Vlan103
    no shutdown
    vrf member RED
    ip address 10.103.103.1/24 tag 2031
    fabric forwarding mode anycast-gateway
!
interface Vlan203
    no shutdown
    vrf member BLUE
    ip address 10.203.203.1/24 tag 2032
    fabric forwarding mode anycast-gateway
!
interface Vlan2031
    description L3 VNI for VRF RED
    no shutdown
    vrf member RED
    ip forward
!
interface Vlan2032
    description L3 VNI for VRF BLUE
    no shutdown
    vrf member BLUE
    ip forward
!
interface nve1
    no shutdown
    source-interface loopback0
    host-reachability protocol bgp
    member vni 10103
        suppress-arp
        mcast-group 239.103.103.103
    member vni 10203
        suppress-arp
        mcast-group 239.203.203.203
    member vni 12031 associate-vrf
    member vni 12032 associate-vrf
!
interface Ethernet1/2
    description SRV6
    switchport access vlan 203
    spanning-tree port type edge
!
interface Ethernet1/5
    no switchport
    medium p2p
    bfd interval 250 min_rx 250 multiplier 4
    no ip redirects
    ip unnumbered loopback1

```

```

ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface Ethernet1/6
no switchport
medium p2p
bfd interval 250 min_rx 250 multiplier 4
no ip redirects
ip unnumbered loopback1
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface Ethernet1/15.55
description UPLINK TO BBR
encapsulation dot1Q 55
vrf member RED
ip address 10.0.55.56/24 tag 2001
!
interface Ethernet1/15.56
description UPLINK TO BBR
encapsulation dot1Q 56
vrf member BLUE
ip address 10.0.56.56/24 tag 2002
!
interface loopback0
description OVERLAY LOOPBACK
ip address 2.2.2.56/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface loopback1
description UNDERLAY LOOPBACK
ip address 1.1.1.56/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
router ospf UNDERLAY
bfd
router-id 1.1.1.56
!
router bgp 65000
router-id 2.2.2.56
address-family l2vpn evpn
template peer SPINE_PEERS
update-source loopback0
address-family l2vpn evpn
send-community both
neighbor 2.2.2.75 remote-as 65000
inherit peer SPINE_PEERS
description 7K5
neighbor 2.2.2.76 remote-as 65000
inherit peer SPINE_PEERS
description 7K6
vrf BLUE
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map RM-BLUE-ROUTES
redistribute static route-map RM-BLUE-ROUTES
default-information originate
vrf RED
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map RM-RED-ROUTES
redistribute static route-map RM-RED-ROUTES
default-information originate
!
evpn
vni 10103 12
rd auto
route-target import auto
route-target export auto
vni 10203 12
rd auto
route-target import auto
route-target export auto
!
ngoam install acl

N7K5:
hostname N7K5
!
nv overlay evpn
feature ospf
feature bgp
feature pim
feature udld
feature bfd
feature nxapi
!
feature ngoam
!
ip pim rp-address 7.7.7.7 group-list 224.0.0.0/4 bidir
ip pim ssm range 232.0.0.0/8
ip pim bfd
!
interface Ethernet1/19
medium p2p
bfd interval 250 min_rx 250 multiplier 4
no ip redirects

```

```

ip unnumbered loopback0
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
no shutdown
!
interface Ethernet1/20
medium p2p
bfd interval 250 min_rx 250 multiplier 4
no ip redirects
ip unnumbered loopback1
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
no shutdown
!
interface loopback0
description OVERLAY LOOPBACK
ip address 2.2.2.75/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface loopback1
description UNDERLAY LOOPBACK
ip address 1.1.1.75/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface loopback7
description Anycast RP Address
ip address 7.7.7.6/31
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
router ospf UNDERLAY
bfd
router-id 1.1.1.75
!
router bgp 65000
router-id 2.2.2.75
address-family l2vpn evpn
template peer LEAF_PEERS
update-source loopback0
address-family l2vpn evpn
send-community both
route-reflector-client
neighbor 2.2.2.0/24 remote-as 65000
inherit peer LEAF_PEERS

N7K6:
hostname N7K6
!
nv overlay evpn
feature ospf
feature bgp
feature pim
feature udld
feature bfd
feature nxapi
!
feature ngoam
!
ip pim rp-address 7.7.7.7 group-list 224.0.0.0/4 bidir
ip pim ssm range 232.0.0.0/8
ip pim bfd
!
interface Ethernet1/19
medium p2p
bfd interval 250 min_rx 250 multiplier 4
no ip redirects
ip unnumbered loopback1
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
no shutdown
!
interface Ethernet1/20
medium p2p
bfd interval 250 min_rx 250 multiplier 4
no ip redirects
ip unnumbered loopback1
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
no shutdown
!
interface loopback0
description OVERLAY LOOPBACK
ip address 2.2.2.76/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface loopback1
description UNDERLAY LOOPBACK
ip address 1.1.1.76/32
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
interface loopback7

```

```
description Anycast RP Address
ip address 7.7.7.5/30
ip ospf network point-to-point
ip router ospf UNDERLAY area 0.0.0.0
ip pim sparse-mode
!
router ospf UNDERLAY
  bfd
  router-id 1.1.1.76
!
router bgp 65000
  router-id 2.2.2.76
  address-family 12vpn evpn
  template peer LEAF_PEERS
  update-source loopback0
  address-family 12vpn evpn
  send-community both
  route-reflector-client
  neighbor 2.2.2.0/24 remote-as 65000
  inherit peer LEAF_PEERS
```



« [VXLAN BGP EVPN Layer 2 Fabric | ACI Fabric Initialization](#) »

[^ back to top](#)

© 2017 INE Inc., All Rights Reserved