

A solid blue vertical bar with rounded ends, positioned to the left of the section header.

IP fabric design guide for Eltex switches

1	Introduction	4
2	Terms and abbreviations.....	5
3	Spine-Leaf architecture	6
4	Installing licenses	8
5	Configuring underlay. IS-IS.....	9
5.1	Configuring Spine.....	10
5.2	Configuring Leaf	13
5.3	Checking underlay configuration	16
5.4	Expected result	17
6	Configuring underlay. OSPF	18
6.1	Configuring Spine.....	18
6.2	Configuring Leaf	21
6.3	Checking underlay configuration	24
6.4	Expected result	25
7	Configuring overlay. VXLAN.....	26
7.1	Configuring VXLAN.....	26
7.2	Checking VXLAN configuration	27
7.3	ARP suppression.....	33
7.4	Expected result	33
8	Configuring overlay. Multicast VXLAN	34
8.1	Configuring Spine.....	34
8.2	Configuring Leaf	35
8.3	Configuring Multicast VXLAN	36
8.4	Checking Multicast VXLAN configuration.....	39
8.5	Expected result	41
9	Configuring overlay. Symmetric IRB	42
9.1	Configuring Symmetric IRB	42
9.2	Checking Symmetric IRB configuration.....	45
9.3	Anycast gateway.....	47
9.4	Gateway-ip for type 5 routes	49
9.5	DHCP relay	50
9.6	Expected result	55

10	EVPN multihoming.....	56
10.1	Configuring EVPN multihoming	57
10.2	Checking EVPN multihoming configuration.....	58
10.3	Expected result	60
11	Management network.....	61
11.1	OOB network	61
11.2	Device configurations.....	62
12	Appendix 1.....	67
12.1	Configurations for IS-IS	68
12.2	Configurations for OSPF.....	78
12.3	Configurations for Multicast VXLAN	88
12.4	Configurations for Symmetric IRB	102
12.5	Configurations for EVPN multihoming	121

1 Introduction

The purpose of this guide is to provide the basic steps necessary to design an IP fabric based on Eltex equipment using EVPN/VXLAN.

The target audience includes network specialists, system integrators, partners, and customers who use or plan to use Eltex equipment.

2 Terms and abbreviations

Anycast gateway — allows the same IP and MAC gateway addresses to be used on all Leaf devices that are members of the same L3VNI.

ARP suppression — allows the Leaf device to respond to an ARP request on behalf of a remote host without forwarding the ARP request via VXLAN.

BFD (Bidirectional Forwarding Detection protocol) — a protocol designed for rapid detection of link failures.

ECMP (Equal-cost multi-path routing) — a load balancing technology that allows packets to be transmitted to a single recipient via multiple 'best routes'. This feature is designed to distribute load and optimize network bandwidth.

ESI (Ethernet Segment Identifier) — a 10-byte Ethernet segment identifier that is unique within an IP fabric.

Ethernet-segment (ES) — a set of links forming a LAG connecting a client device to a group of Leaves.

EVPN (Ethernet Virtual Private Network) — a standard defined in RFC 7432. EVPN is an extension of the BGP protocol (address-family, AFI: 25, SAFI: 70) that functions as a management tool for creating L2VPN and L3VPN services. The extension allows the network to transmit information about end devices, such as MAC and IP addresses.

IBGP (Internal BGP) — used to connect BGP neighbors within a single autonomous system.

IGP (Interior Gateway Protocol) — used to transmit information about routes within an autonomous system. These include IS-IS and OSPF protocols.

Ingress replication — a VXLAN mode in which BUM traffic replication is performed on the incoming VTEP.

IP fabric — a network infrastructure based on the IP protocol that allows creating multiple symmetrical paths between all devices in the IP fabric.

L3VNI — a VNI used for routing.

Leaf — an access layer device of an IP fabric.

MSDP (Multicast Source Discovery Protocol) — used to exchange information about multicast traffic sources between Spines.

Multicast VXLAN — a VXLAN mode in which BUM traffic replication is performed via PIM.

OOB interface — a device port for remote management. Management is carried out over the network separately from the data transmission channel.

Overlay network — a logical network created over another underlay network and using its infrastructure as transport.

PIM (Protocol Independent Multicast) — a multicast routing protocol for IP networks.

POD (Point Of Delivery) — a separate group of devices in the Clos topology (first-level Spines and Leaves), where the Spine has connections only to the Leaves of its group and has no connection to the Leaves of the rest of the IP fabric.

Route target (RT) — extended BGP community.

Spine — a central device in an IP fabric that has connections to all Leaves (to all Leaves of its POD if there is a POD on the network).

Underlay network — a basic physical network providing connectivity between all devices.


VNI (Virtual Network Identifier) — a 24-bit network identifier within VXLAN.

VTEP (Virtual Tunnel End Point) — device where a VXLAN tunnel begins or ends. Traffic is encapsulated in a VXLAN at the source VTEP and decapsulated at the remote VTEP.

VXLAN (Virtual eXtensible Local Area Network) — a RFC 7348 standard technology for creating a virtual (overlay) network over existing IP infrastructure.

3 Spine-Leaf architecture

An example topology of an IP fabric created according to the Spine-Leaf architecture is given below.




Such a topology is also known as Folded Clos.

Terminal equipment is connected to Leaf devices that act as access switches. Leaf devices, in turn, are connected to Spine (switching-layer) devices.

The purpose of this layer is to route packets from one Leaf to another as quickly as possible. Thus, between each pair of Leaves, there are a number of equivalent paths equal to the number of Spines.

In the diagram below, the IP fabric used as an example has been expanded twice.



The IP fabric is divided into two PODs (see [Terms and abbreviations](#)), and an additional switching layer (SuperSpine or Second-layer Spine) is implemented. The diagram shows that the PODs are isolated from each other: the First-layer Spine (Spine_11-22 in the diagram) only has connections to a specific group of Leaf devices that belong to its POD. PODs are connected to each other via the Second-layer Spine.

The diagram uses the following protocols and technologies:

- Due to the use of EVPN, the main routing protocol is BGP. BGP sessions are established between the loopback interfaces of devices. All devices have a single autonomous system number, forming an iBGP space.
- L2VPN EVPN address-family is used in the BGP protocol settings. Message exchange within this address-family ensures the operation of EVPN technology.
- IS-IS and OSPF protocols are used as IGP (Internal Gateway Protocol) in this guide. The main task of IGP is to ensure IP connectivity between all loopback interfaces of the network, as BGP sessions and VXLAN tunnels are established between them.
- For quick detection of link failures, the BFD protocol is used. It can detect a link failure in less than 1 second. The minimum configurable response time is 150 ms.
- BGP route reflector (RR) is a BGP configuration that allows a router to act as a route reflector. Since Clos topology eliminates full connectivity and uses iBGP, route information distribution is limited to direct links only. RR allows spine-level switches to retransmit routing information updates between Leaf switches.
- The use of ECMP technology in the Clos topology is essential to eliminate unused links. The load is evenly balanced on all links available between any pair of Leaf devices. The fault tolerance of the network is increased.
- Jumbo-frame – transmission of large frames of up to 10,200 bytes is supported.
- Spanning Tree protocols are disabled on all IP fabric devices.

4 Installing licenses

Support for the BGP protocol and EVPN/VXLAN technology is provided under license. Ensure that the appropriate licences are obtained. If licences are not present, install them.

1. The **show license** command output if no licenses are installed:

```
console#show license
Features installed:

      Feature      Licenses  Licenses
      installed  used      Active
-----
Licenses installed:
```

2. To install the license, download it to the device using the **boot license source_url** command. Example:

```
console#boot license tftp://192.168.1.1/licensefile

console#boot license tftp://192.168.1.1/licensefile
02-Jun-2022 12:01:49 %COPY-I-FILECPY: Files Copy - source URL tftp://192.168.1.1/licensefile
destination URL flash://system/licenses/licensefile
02-Jun-2022 12:01:49 %LICENSE-I-INSTALL: License file licensefile was installed
02-Jun-2022 12:01:49 %COPY-N-TRAP: The copy operation was completed successfully

Copy: 1181 bytes copied in 00:00:01 [hh:mm:ss]
```

3. Reload the device after downloading the file to apply the license.

4. The **show license** command output if the licenses required are installed:

```
console#show license
Features installed:


      Feature      Licenses  Licenses
      installed  used      Active
-----
BGP             1           1       Yes
EVPN            1           1       Yes

Licenses installed:

License name: licensefile
License version: 1.0
Valid for device: ES7A000000 (e0:d9:e3:00:00:00)
Status: Active
Features:
  BGP, EVPN
```


5 Configuring underlay. IS-IS

The network topology is shown below.



⚠ IP addressing in this guide is provided for illustrative purposes only.

Addresses of loopback interfaces (labelled on the diagram):

Spine_1: 10.0.1.1

Spine_2: 10.0.2.2

Leaf_1: 10.0.0.1

Leaf_2: 10.0.0.2

Leaf_3: 10.0.0.3

Linking networks:

Spine_1-Leaf_1: 172.16.1.0/30

Spine_1-Leaf_2: 172.16.2.0/30

Spine_1-Leaf_3: 172.16.3.0/30

Spine_2-Leaf_1: 172.16.1.4/30

Spine_2-Leaf_2: 172.16.2.4/30

Spine_2-Leaf_3: 172.16.3.4/30


5.1 Configuring Spine

Perform the initial configuration of the switch:

```
console(config)#no spanning-tree
console(config)#port jumbo-frame
This setting will take effect only after copying running configuration to startup configuration
and resetting the device
console(config)#ip maximum-paths 32
Warning! New value will be applied only after reboot
console(config)#hostname Spine_1
Spine_1(config)#
```

Where:

- `no spanning-tree` – disabling STP;
- `port jumbo-frame` – enabling support for jumbo frames;
- `ip maximum-paths 32` – setting the maximum number of paths that can be established in FIB for each route, thereby enabling ECMP;
- `hostname Spine_1` – setting the device name.

 The **port jumbo-frame** and **ip maximum-paths 32** settings take effect only after the device is reloaded. To do this, save the configuration and reload:

```
Spine_1#write
Overwrite file [startup-config]... (Y/N)[N] ?Y
23-Jun-2022 07:13:16 %COPY-I-FILECOPY: Files Copy - source URL running-config destination URL
flash://system/configuration/startup-config
23-Jun-2022 07:13:16 %COPY-N-TRAP: The copy operation was completed successfully
Copy succeeded
Spine_1#reload
This command will reset the whole system and disconnect your current session. Do you want to
continue ? (Y/N)[N] Y
Shutting down ...
```

The application of settings after reloading can be checked in the output of the following show commands. Example:

```
Spine_1#show ports jumbo-frame

Jumbo frames are enabled
Jumbo frames will be enabled after reset

Spine_1#show ip route
Maximum Parallel Paths: 32 (32 after reset)
Load balancing: src-dst-mac-ip
IP Forwarding: enabled
Codes: > - best, C - connected, S - static,
        R - RIP,
        O - OSPF intra-area, OIA - OSPF inter-area,
        OE1 - OSPF external 1, OE2 - OSPF external 2,
        B - BGP, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, ia - IS-IS inter area
```

The **Jumbo frames are enabled** and **Maximum Parallel Paths: 32 (32 after reset)** lines indicate that the respective settings have been successfully enabled.

Configure the interfaces.

To simplify configuration via CLI, first use the **terminal no prompt** feature, which disables the need for confirmation before executing certain commands:

```
Spine_1#terminal no prompt
```

Configuring interfaces:

```
Spine_1(config)#interface HundredGigabitEthernet1/0/1
Spine_1(config-if)# description Leaf_1
Spine_1(config-if)# ip address 172.16.1.2 255.255.255.252
Spine_1(config-if)# ip router isis
Spine_1(config-if)# isis network point-to-point
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface HundredGigabitEthernet1/0/2
Spine_1(config-if)# description Leaf_2
Spine_1(config-if)# ip address 172.16.2.2 255.255.255.252
Spine_1(config-if)# ip router isis
Spine_1(config-if)# isis network point-to-point
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface HundredGigabitEthernet1/0/3
Spine_1(config-if)# description Leaf_3
Spine_1(config-if)# ip address 172.16.3.2 255.255.255.252
Spine_1(config-if)# ip router isis
Spine_1(config-if)# isis network point-to-point
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface loopback1
Spine_1(config-if)# ip address 10.0.1.1 255.255.255.255
Spine_1(config-if)#exit
```

Enabling the IS-IS protocol:

```
Spine_1(config)#router isis
Spine_1(router-isis)# address-family ipv4 unicast
Spine_1(router-isis-af)# redistribute connected
Spine_1(router-isis-af)# exit
Spine_1(router-isis)# net 49.0001.1111.1111.1111.00
Spine_1(router-isis)#exit
```

Enabling the BGP protocol:

 Support for the BGP protocol is provided under license (see [Installing licences](#)).

```
Spine_1(config)#router bgp 65500
Spine_1(router-bgp)# bgp router-id 10.0.1.1
This action will reset all neighbor connections and clear BGP routing table.
Spine_1(router-bgp)# address-family ipv4 unicast
Spine_1(router-bgp-af)# exit
Spine_1(router-bgp)# address-family l2vpn evpn
This action will reset all neighbor connections and clear BGP routing table.
Spine_1(router-bgp-af)# exit
Spine_1(router-bgp)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbrgrp)# remote-as 65500
This action will reset connection with all neighbors in peer group.
Spine_1(router-bgp-nbrgrp)# update-source loopback 1
This action will reset connection with all neighbors in peer group.
Spine_1(router-bgp-nbrgrp)# fall-over bfd
Spine_1(router-bgp-nbrgrp)# route-reflector-client
Spine_1(router-bgp-nbrgrp)# exit
Spine_1(router-bgp)# neighbor 10.0.0.1
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)# neighbor 10.0.0.2
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)# neighbor 10.0.0.3
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)#exit
```

Other Spine devices are configured in the same way, with the necessary changes made according to the diagram and IP addressing plan.

The following parameters are subject to change from device to device:

- Hostname of the device;
- Description on interfaces;
- IP addresses of interfaces (physical + loopback);
- IS-IS net identifier;
- BGP router-id;
- IP addresses of BGP neighbors.


5.2 Configuring Leaf

Perform the initial configuration of the switch:

```
console(config)#no spanning-tree
console(config)#port jumbo-frame
This setting will take effect only after copying running configuration to startup configuration
and resetting the device
console(config)#ip maximum-paths 32
Warning! New value will be applied only after reboot
console(config)#hostname Leaf_1
Leaf_1(config)#
```

Where:

- no spanning-tree – disabling STP;
- port jumbo-frame – enabling support for jumbo frames;
- ip maximum-paths 32 – setting the maximum number of paths that can be established in FIB for each route, thereby enabling ECMP;
- hostname Leaf_1 – setting the device name.

 The **port jumbo-frame** and **ip maximum-paths 32** settings take effect only after the device is reloaded. To do this, save the configuration and reload:

```
Leaf_1#write
Overwrite file [startup-config].... (Y/N)[N] ?Y
23-Jun-2022 07:13:16 %COPY-I-FILECPY: Files Copy - source URL running-config destination URL
flash://system/configuration/startup-config
23-Jun-2022 07:13:16 %COPY-N-TRAP: The copy operation was completed successfully
Copy succeeded
Leaf_1#reload
This command will reset the whole system and disconnect your current session. Do you want to
continue ? (Y/N)[N] Y
Shutting down ...
```

The application of settings after reloading can be checked in the output of the following show commands. Example:

```
Leaf_1#show ports jumbo-frame

Jumbo frames are enabled
Jumbo frames will be enabled after reset

Leaf_1#show ip route
Maximum Parallel Paths: 32 (32 after reset)
Load balancing: src-dst-mac-ip
IP Forwarding: enabled
Codes: > - best, C - connected, S - static,
        R - RIP,
        O - OSPF intra-area, OIA - OSPF inter-area,
        OE1 - OSPF external 1, OE2 - OSPF external 2,
        B - BGP, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, ia - IS-IS inter area
```

The **Jumbo frames are enabled** and **Maximum Parallel Paths: 32 (32 after reset)** lines indicate that the respective settings have been successfully enabled.

Configure the interfaces.

To simplify configuration via CLI, you can first use the **terminal no prompt** feature, which disables the need for confirmation before executing certain commands:

```
Leaf_1#terminal no prompt
```

Configuring interfaces:

```
Leaf_1(config)#interface HundredGigabitEthernet1/0/1
Leaf_1(config-if)# description Spine_1
Leaf_1(config-if)# ip address 172.16.1.1 255.255.255.252
Leaf_1(config-if)# ip router isis
Leaf_1(config-if)# isis network point-to-point
This action will reset all neighbor connections on the interface.
Leaf_1(config-if)#exit
Leaf_1(config)#interface HundredGigabitEthernet1/0/2
Leaf_1(config-if)# description Spine_2
Leaf_1(config-if)# ip address 172.16.1.5 255.255.255.252
Leaf_1(config-if)# ip router isis
Leaf_1(config-if)# isis network point-to-point
This action will reset all neighbor connections on the interface.
Leaf_1(config-if)#exit
Leaf_1(config)#interface loopback1
Leaf_1(config-if)# ip address 10.0.0.1 255.255.255.255
Leaf_1(config-if)#exit
```

Enabling the IS-IS protocol:

```
Leaf_1(config)#router isis
Leaf_1(router-isis)# address-family ipv4 unicast
Leaf_1(router-isis-af)# redistribute connected
Leaf_1(router-isis-af)# exit
Leaf_1(router-isis)# net 49.0001.0001.0001.0001.00
Leaf_1(router-isis)#exit
```

Enabling the BGP protocol:

 Support for the BGP protocol is provided under license (see [Installing licenses](#)).

```
Leaf_1(config)#router bgp 65500
Leaf_1(router-bgp)# bgp router-id 10.0.0.1
This action will reset all neighbor connections and clear BGP routing table.
Leaf_1(router-bgp)# address-family ipv4 unicast
Leaf_1(router-bgp-af)# exit
Leaf_1(router-bgp)# address-family l2vpn evpn
This action will reset all neighbor connections and clear BGP routing table.
Leaf_1(router-bgp-af)# exit
Leaf_1(router-bgp)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbrgrp)# remote-as 65500
This action will reset connection with all neighbors in peer group.
Leaf_1(router-bgp-nbrgrp)# update-source loopback 1
This action will reset connection with all neighbors in peer group.
Leaf_1(router-bgp-nbrgrp)# fall-over bfd
Leaf_1(router-bgp-nbrgrp)# exit
Leaf_1(router-bgp)# neighbor 10.0.1.1
Leaf_1(router-bgp-nbr)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbr)# address-family ipv4 unicast
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# exit
Leaf_1(router-bgp)# neighbor 10.0.2.2
Leaf_1(router-bgp-nbr)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbr)# address-family ipv4 unicast
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# exit
Leaf_1(router-bgp)#exit
```

Other Leaf devices are configured in the same way, with the necessary changes made according to the diagram and IP addressing plan.

The following parameters are subject to change from device to device:

- Hostname of the device;
- Description on interfaces;
- IP addresses of interfaces (physical + loopback);
- IS-IS net identifier;
- BGP router-id;
- IP addresses of BGP neighbors.

5.3 Checking underlay configuration

After completing the above settings, verify the IS-IS, BGP, and BFD neighborhoods.

The commands used:

```
show isis neighbors
show ip bgp neighbors
show ip bfd neighbors
```

Example of the above show commands execution on Spine_1:

```
Spine_1#show isis neighbors
```

System Id	Interface	Type	SNPA	State	Holdtime (s)	Circuit Id
Leaf_1	hu1/0/1	L1-2	e0d9.e326.d600	Up	29	
Leaf_2	hu1/0/2	L1-2	e0d9.e3f8.6e00	Up	29	
Leaf_3	hu1/0/3	L1-2	e0d9.e3d7.ea80	Up	27	

```
Spine_1#show ip bgp neighbors
```

BGP neighbor	Remote AS	Router ID	State	Uptime	Hold Time	Keepalive
10.0.0.1	65500	10.0.0.1	ESTABLISHED	00,00:01:15	90	30
10.0.0.2	65500	10.0.0.2	ESTABLISHED	00,00:00:58	90	30
10.0.0.3	65500	10.0.0.3	ESTABLISHED	00,00:00:44	90	30

```
Spine_1#show ip bfd neighbors
```

Neighbor	Local	State	Last Down Diag
10.0.0.1	10.0.1.1	Up	No Diagnostic
10.0.0.2	10.0.1.1	Up	No Diagnostic
10.0.0.3	10.0.1.1	Up	No Diagnostic

Example of the above show commands execution on Leaf_1:

```
Leaf_1#show isis neighbors
```

System Id	Interface	Type	SNPA	State	Holdtime (s)	Circuit Id
Spine_1	hu1/0/1	L1-2	cc9d.a253.d680	Up	29	
Spine_2	hu1/0/2	L1-2	e0d9.e317.6b40	Up	28	

```
Leaf_11#show ip bgp neighbors
```

BGP neighbor Keepalive	Remote AS	Router ID	State	Uptime	Hold Time
10.0.1.1	65500	10.0.1.1	ESTABLISHED	00,00:01:25	90
10.0.2.2	65500	10.0.2.2	ESTABLISHED	00,00:01:23	90

```
Leaf_11#show ip bfd neighbors
```

Neighbor	Local	State	Last Down Diag
10.0.1.1	10.0.0.1	Up	No Diagnostic
10.0.2.2	10.0.0.1	Up	No Diagnostic

If all neighborhoods are successfully established, IP connectivity between all devices must be ensured. To check this, run the icmp ping command from any device to any other device, with IP addresses of device loopback interfaces specified as src and dst.

Example of ping from Leaf_1 to Leaf_3:

```
Leaf_1#ping 10.0.0.3 source 10.0.0.1
```

```
Pinging 10.0.0.3 with 18 bytes of data:
```

```
18 bytes from 10.0.0.3: icmp_seq=1. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=2. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=3. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=4. time=0 ms
```

```
----10.0.0.3 PING Statistics----
```

```
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms) min/avg/max = 0/0/0
```

5.4 Expected result

- All IS-IS and BFD neighborhoods are in the UP state. All BGP neighborhoods are in the ESTABLISHED state.
- IP connectivity between all the devices is ensured.

6 Configuring underlay. OSPF

The network topology is similar to the topology given in [Configuring underlay. IS-IS](#).


6.1 Configuring Spine

Perform the initial configuration of the switch:

```
console(config)#no spanning-tree
console(config)#port jumbo-frame
This setting will take effect only after copying running configuration to startup configuration
and resetting the device
console(config)#ip maximum-paths 32
Warning! New value will be applied only after reboot
console(config)#hostname Spine_1
Spine_1(config)#
```

Where:

- `no spanning-tree` – disabling STP;
- `port jumbo-frame` – enabling support for jumbo frames;
- `ip maximum-paths 32` – setting the maximum number of paths that can be established in FIB for each route, thereby enabling ECMP;
- `hostname Spine_1` – setting the device name.

 The **port jumbo-frame** and **ip maximum-paths 32** settings take effect only after the device is reloaded. To do this, save the configuration and reload:

```
Spine_1#write
Overwrite file [startup-config].... (Y/N)[N] ?Y
23-Jun-2022 07:13:16 %COPY-I-FILECPY: Files Copy - source URL running-config destination URL
flash://system/configuration/startup-config
23-Jun-2022 07:13:16 %COPY-N-TRAP: The copy operation was completed successfully
Copy succeeded
Spine_1#reload
This command will reset the whole system and disconnect your current session. Do you want to
continue ? (Y/N)[N] Y
Shutting down ...
```

The application of settings after reloading can be checked in the output of the following show commands. Example:

```
Spine_1#show ports jumbo-frame

Jumbo frames are enabled
Jumbo frames will be enabled after reset


Spine_1#show ip route
Maximum Parallel Paths: 32 (32 after reset)
Load balancing: src-dst-mac-ip
IP Forwarding: enabled
Codes: > - best, C - connected, S - static,
        R - RIP,
        O - OSPF intra-area, OIA - OSPF inter-area,
        OE1 - OSPF external 1, OE2 - OSPF external 2,
        B - BGP, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, ia - IS-IS inter area
```

The **Jumbo frames are enabled** and **Maximum Parallel Paths: 32 (32 after reset)** lines indicate that the respective settings have been successfully enabled.

Configure the interfaces.


To simplify configuration via CLI, first use the **terminal no prompt** feature, which disables the need for confirmation before executing certain commands:

```
Spine_1#terminal no prompt
```

Configuring interfaces:

```
Spine_1(config)#interface HundredGigabitEthernet1/0/1
Spine_1(config-if)# description Leaf_1
Spine_1(config-if)# ip address 172.16.1.2 255.255.255.252
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface HundredGigabitEthernet1/0/2
Spine_1(config-if)# description Leaf_2
Spine_1(config-if)# ip address 172.16.2.2 255.255.255.252
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface HundredGigabitEthernet1/0/3
Spine_1(config-if)# description Leaf_3
Spine_1(config-if)# ip address 172.16.3.2 255.255.255.252
This action will reset all neighbor connections on the interface.
Spine_1(config-if)#exit
Spine_1(config)#interface loopback1
Spine_1(config-if)# ip address 10.0.1.1 255.255.255.255
Spine_1(config-if)#exit
```

Enabling the OSPF protocol:

 In order to minimize traffic interruption when changing the state of Spine-Leaf links, it is recommended to use the OSPF **timers spf delay 0** setting. It allows removing the delay before the next SPF calculation and, accordingly, removes the delay before updating the routing table when changes occur in the OSPF topology.

```
Spine_1(config)#router ospf 1
Spine_1(router_ospf_process)# network 172.16.1.2 area 0.0.0.0
Spine_1(router_ospf_process)# network 172.16.2.2 area 0.0.0.0
Spine_1(router_ospf_process)# network 172.16.3.2 area 0.0.0.0
Spine_1(router_ospf_process)# router-id 10.0.1.1
Spine_1(router_ospf_process)# timers spf delay 0
Spine_1(router_ospf_process)# redistribute connected subnets
Spine_1(router_ospf_process)#exit
Spine_1(config)#interface ip 172.16.1.2
Spine_1(config-if)# ip ospf network point-to-point
Spine_1(config-if)#exit
Spine_1(config)#interface ip 172.16.2.2
Spine_1(config-if)# ip ospf network point-to-point
Spine_1(config-if)#exit
Spine_1(config)#interface ip 172.16.3.2
Spine_1(config-if)# ip ospf network point-to-point
Spine_1(config-if)#exit
```

Enabling the BGP protocol:

 Support for the BGP protocol is provided under license (see [Installing licenses](#)).

```
Spine_1(config)#router bgp 65500
Spine_1(router-bgp)# bgp router-id 10.0.1.1
This action will reset all neighbor connections and clear BGP routing table.
Spine_1(router-bgp)# address-family ipv4 unicast
Spine_1(router-bgp-af)# exit
Spine_1(router-bgp)# address-family l2vpn evpn
This action will reset all neighbor connections and clear BGP routing table.
Spine_1(router-bgp-af)# exit
Spine_1(router-bgp)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbrgrp)# remote-as 65500
This action will reset connection with all neighbors in peer group.
Spine_1(router-bgp-nbrgrp)# update-source loopback 1
This action will reset connection with all neighbors in peer group.
Spine_1(router-bgp-nbrgrp)# fall-over bfd
Spine_1(router-bgp-nbrgrp)# route-reflector-client
Spine_1(router-bgp-nbrgrp)# exit
Spine_1(router-bgp)# neighbor 10.0.0.1
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)# neighbor 10.0.0.2
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)# neighbor 10.0.0.3
Spine_1(router-bgp-nbr)# peer-group LEAF_GROUP
Spine_1(router-bgp-nbr)# address-family ipv4 unicast
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Spine_1(router-bgp-nbr-af)# exit
Spine_1(router-bgp-nbr)# exit
Spine_1(router-bgp)#exit
```

Other Spine devices are configured in the same way, with the necessary changes made according to the diagram and IP addressing plan.

The following parameters are subject to change from device to device:

- Hostname of the device;
- Description on interfaces;
- IP addresses of interfaces (physical + loopback);
- IP interfaces on which OSPF is enabled;
- OSPF router-id;
- BGP router-id;
- IP addresses of BGP neighbors.


6.2 Configuring Leaf

Perform the initial configuration of the switch:

```
console(config)#no spanning-tree
console(config)#port jumbo-frame
This setting will take effect only after copying running configuration to startup configuration
and resetting the device
console(config)#ip maximum-paths 32
Warning! New value will be applied only after reboot
console(config)#hostname Leaf_1
Leaf_1(config)#
```

Where:

- no spanning-tree – disabling STP;
- port jumbo-frame – enabling support for jumbo frames;
- ip maximum-paths 32 – setting the maximum number of paths that can be established in FIB for each route, thereby enabling ECMP;
- hostname Leaf_1 – setting the device name.

 The **port jumbo-frame** and **ip maximum-paths 32** settings take effect only after the device is reloaded. To do this, save the configuration and reload:

```
Leaf_1#write
Overwrite file [startup-config]... (Y/N)[N] ?Y
23-Jun-2022 07:13:16 %COPY-I-FILECOPY: Files Copy - source URL running-config destination URL
flash://system/configuration/startup-config
23-Jun-2022 07:13:16 %COPY-N-TRAP: The copy operation was completed successfully
Copy succeeded
Leaf_1#reload
This command will reset the whole system and disconnect your current session. Do you want to
continue ? (Y/N)[N] Y
Shutting down ...
```

The application of settings after reloading can be checked in the output of the following show commands. Example:

```
Leaf_1#show ports jumbo-frame

Jumbo frames are enabled
Jumbo frames will be enabled after reset

Leaf_1#show ip route
Maximum Parallel Paths: 32 (32 after reset)
Load balancing: src-dst-mac-ip
IP Forwarding: enabled
Codes: > - best, C - connected, S - static,
        R - RIP,
        O - OSPF intra-area, OIA - OSPF inter-area,
        OE1 - OSPF external 1, OE2 - OSPF external 2,
        B - BGP, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, ia - IS-IS inter area
```

The **Jumbo frames are enabled** and **Maximum Parallel Paths: 32 (32 after reset)** lines indicate that the respective settings have been successfully enabled.

Configure the interfaces.


To simplify configuration via CLI, you can first use the **terminal no prompt** feature, which disables the need for confirmation before executing certain commands:

```
Leaf_1#terminal no prompt
```

Configuring interfaces:

```
Leaf_1(config)#interface HundredGigabitEthernet1/0/1
Leaf_1(config-if)# description Spine_1
Leaf_1(config-if)# ip address 172.16.1.1 255.255.255.252
This action will reset all neighbor connections on the interface.
Leaf_1(config-if)#exit
Leaf_1(config)#interface HundredGigabitEthernet1/0/2
Leaf_1(config-if)# description Spine_2
Leaf_1(config-if)# ip address 172.16.1.5 255.255.255.252
This action will reset all neighbor connections on the interface.
Leaf_1(config-if)#exit
Leaf_1(config)#interface loopback1
Leaf_1(config-if)# ip address 10.0.0.1 255.255.255.255
Leaf_1(config-if)#exit
```

Enabling the OSPF protocol:

 In order to minimize traffic interruption when changing the state of Spine-Leaf links, it is recommended to use the OSPF **timers spf delay 0** setting. It allows removing the delay before the next SPF calculation and, accordingly, removes the delay before updating the routing table when changes occur in the OSPF topology.

```
Leaf_1(config)#router ospf 1
Leaf_1(router_ospf_process)#network 172.16.1.1 area 0.0.0.0
Leaf_1(router_ospf_process)#network 172.16.1.5 area 0.0.0.0
Leaf_1(router_ospf_process)#router-id 10.0.0.1
Leaf_1(router_ospf_process)#timers spf delay 0
Leaf_1(router_ospf_process)#redistribute connected subnets
Leaf_1(router_ospf_process)#exit
Leaf_1(config)#interface ip 172.16.1.1
Leaf_1(config-if)#ip ospf network point-to-point
Leaf_1(config-if)#exit
Leaf_1(config)#interface ip 172.16.1.5
Leaf_1(config-if)#ip ospf network point-to-point
Leaf_1(config-if)#exit
```

Enabling the BGP protocol:

 Support for the BGP protocol is provided under license (see [Installing licenses](#)).

```
Leaf_1(config)#router bgp 65500
Leaf_1(router-bgp)# bgp router-id 10.0.0.1
This action will reset all neighbor connections and clear BGP routing table.
Leaf_1(router-bgp)# address-family ipv4 unicast
Leaf_1(router-bgp-af)# exit
Leaf_1(router-bgp)# address-family l2vpn evpn
This action will reset all neighbor connections and clear BGP routing table.
Leaf_1(router-bgp-af)# exit
Leaf_1(router-bgp)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbrgrp)# remote-as 65500
This action will reset connection with all neighbors in peer group.
Leaf_1(router-bgp-nbrgrp)# update-source loopback 1
This action will reset connection with all neighbors in peer group.
Leaf_1(router-bgp-nbrgrp)# fall-over bfd
Leaf_1(router-bgp-nbrgrp)# exit
Leaf_1(router-bgp)# neighbor 10.0.1.1
Leaf_1(router-bgp-nbr)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbr)# address-family ipv4 unicast
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# exit
Leaf_1(router-bgp)# neighbor 10.0.2.2
Leaf_1(router-bgp-nbr)# peer-group SPINE_GROUP
Leaf_1(router-bgp-nbr)# address-family ipv4 unicast
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# address-family l2vpn evpn
This action will reset connection with the neighbor.
Leaf_1(router-bgp-nbr-af)# exit
Leaf_1(router-bgp-nbr)# exit
Leaf_1(router-bgp)#exit
```

Other Leaf devices are configured in the same way, with the necessary changes made according to the diagram and IP addressing plan.

The following parameters are subject to change from device to device:

- Hostname of the device;
- Description on interfaces;
- IP addresses of interfaces (physical + loopback);
- IP interfaces on which OSPF is enabled;
- OSPF router-id;
- BGP router-id;
- IP addresses of BGP neighbors.

6.3 Checking underlay configuration

After completing the above settings, verify the OSPF, BGP, and BFD neighborhoods.

The commands used:

```
show ip ospf neighbor
show ip bgp neighbors
show ip bfd neighbors
```

Example of the above show commands execution on Spine_1:

```
Spine_1#show ip ospf neighbor
```

Neighbor Addr	Neighbor ID	PID	IP Interface	Pri	State	Dead time	Interface
172.16.1.1	10.0.0.1	1	172.16.1.2	1	full/ -	00:00:36	hu1/0/1
172.16.2.1	10.0.0.2	1	172.16.2.2	1	full/ -	00:00:34	hu1/0/2
172.16.3.1	10.0.0.3	1	172.16.3.2	1	full/ -	00:00:37	hu1/0/3

```
Spine_1#show ip bgp neighbors
```

BGP neighbor	Remote AS	Router ID	State	Uptime	Hold Time	Keepalive
10.0.0.1	65500	10.0.0.1	ESTABLISHED	00,00:09:03	90	30
10.0.0.2	65500	10.0.0.2	ESTABLISHED	00,00:02:52	90	30
10.0.0.3	65500	10.0.0.3	ESTABLISHED	00,00:01:45	90	30

```
Spine_1#show ip bfd neighbors
```

Neighbor	Local	State	Last Down Diag
10.0.0.1	10.0.1.1	Up	No Diagnostic
10.0.0.2	10.0.1.1	Up	No Diagnostic
10.0.0.3	10.0.1.1	Up	No Diagnostic

Example of the above show commands execution on Leaf_11:

```
Leaf_1#show ip ospf neighbor
```

Neighbor Addr	Neighbor ID	PID	IP Interface	Pri	State	Dead time	Interface
172.16.1.2	10.0.1.1	1	172.16.1.1	1	full/ -	00:00:37	hu1/0/1
172.16.1.6	10.0.2.2	1	172.16.1.5	1	full/ -	00:00:34	hu1/0/2

```
Leaf_1#show ip bgp neighbors
```

BGP neighbor	Remote AS	Router ID	State	Uptime	Hold Time	Keepalive
10.0.1.1	65500	10.0.1.1	ESTABLISHED	00,00:09:28	90	30
10.0.2.2	65500	10.0.2.2	ESTABLISHED	00,00:05:11	90	30

```
Leaf_1#show ip bfd neighbors
```

Neighbor	Local	State	Last Down Diag
10.0.1.1	10.0.0.1	Up	No Diagnostic
10.0.2.2	10.0.0.1	Up	No Diagnostic

If all neighborhoods are successfully established, IP connectivity between all devices must be ensured. To check this, run the icmp ping command from any device to any other device, with IP addresses of device loopback interfaces specified as src and dst.

Example of ping from Leaf_1 to Leaf_3:

```
Leaf_1#ping 10.0.0.3 source 10.0.0.1
Pinging 10.0.0.3 with 18 bytes of data:

18 bytes from 10.0.0.3: icmp_seq=1. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=2. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=3. time=0 ms
18 bytes from 10.0.0.3: icmp_seq=4. time=0 ms

----10.0.0.3 PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms) min/avg/max = 0/0/0
```

6.4 Expected result

- All OSPF and BFD neighborhoods are in the UP state. All BGP neighborhoods are in the ESTABLISHED state;
- IP connectivity between all the devices is ensured.

7 Configuring overlay. VXLAN

The network topology is similar to the topology given in [Configuring underlay. IS-IS](#) and [Configuring underlay. OSPF](#). Before configuring VXLAN, configure the devices according to one of the sections above.

7.1 Configuring VXLAN

 Support for VXLAN is provided under license (see [Installing licenses](#)).

On Leaf_1, create a VXLAN instance with the name test_vxlan. Set VNI to 101000 and assign VLAN 1000. VLAN 1000 must be created in advance and present in the VLAN database.

```
Leaf_1(config)#vlan database
Leaf_1(config-vlan)#vlan 1000
Leaf_1(config-vlan)#vxlan test_vxlan
Leaf_1(config-vxlan)#vni 101000
Leaf_1(config-vxlan)#vlan 1000
```

By default, the created VXLAN instance is in the **no shutdown** state, i.e., enabled. In the context of VXLAN configuration, it can be disabled with the **shutdown** command. Example:


```
Leaf_1(config-vxlan)#shutdown
```


Create the same VXLAN on Leaf_2 and Leaf_3 devices.

In this example, VLAN 1000 is the client VLAN. The interfaces to Host1 and Host2 must be made members of this VLAN:

```
Leaf_1(config)#interface TenGigabitEthernet1/0/11
Leaf_1(config-if)#description Host1
Leaf_1(config-if)#switchport access vlan 1000

Leaf_3(config)#interface TenGigabitEthernet1/0/11
Leaf_3(config-if)#description Host2
Leaf_3(config-if)#switchport access vlan 1000
```

 In this example, a VXLAN instance is created on all Leaf devices for training purposes to improve the clarity and informativeness of show command outputs. In a real IP fabric, VXLANs are created as needed and on specific Leaf devices.

 When using autonomous system numbers < 65535, it is possible to assign different VLANs to the same VXLAN on different VTEPs without additional settings. In the case of autonomous system numbers > 65535, apply RT (**route-target setting**) to use different VLANs. Its description is given below.

Using an extended community route-target allows different VLANs to be used on different VTEPs for AS >65535.

Example of VXLAN configuration for three Leaf devices using RT and different VLAN numbers (AS > 65535):

```
Leaf_1:

vxlan test_vxlan
vni 101000
vlan 1000
route-target both 65600:100
route-target both 65600:200
exit

Leaf_2:

vxlan test_vxlan
vni 101000
vlan 500
route-target both 65600:100
exit

Leaf_3:

vxlan test_vxlan
vni 101000
vlan 200
route-target both 65600:200
exit
```

All three Leaves use different VLAN numbers. Due to specifying RT, Leaf_1 will establish bidirectional VXLAN tunnels to Leaf_2 and Leaf_3.

However, no tunnel will be established between Leaf_2 and Leaf_3 due to the difference in VLANs assigned to VXLAN and the difference in RT.

For all the Leaves in test_vxlan to be connected to each other, the same RT must be used on all Leaves.

 The maximum number of assigned RTs in a single VXLAN instance is two.

7.2 Checking VXLAN configuration

VXLAN creation can be checked in the following ways:

- in special show commands;
- in the BGP protocol information;
- in the current configuration output.

7.2.1 Show commands

To view information about created VXLANs, use the following commands:

- `show vxlan` — show a table containing information on all created VXLAN instances;
- `show vxlan WORD<1-64>` — show detailed information on a specific VXLAN;
- `show vxlan tunnels` — show all established VXLAN tunnels;
- `show vxlan tunnels WORD<1-64>` — show all established VXLAN tunnels for a specific VXLAN.

Examples of the above commands execution are given below:

```
Leaf_1#show vxlan
```

Name	VNI	VLAN ID	Status	BUM Forwarding	Route Distinguisher
			admin/oper		
test_vxlan	101000	1000	UP/UP	Ingress Repl.	10.0.0.1:1000

```
Leaf_1#show vxlan test_vxlan
```

```
test_vxlan
```

```
VxLAN Network ID is 101000, VLAN ID is 1000
```

```
Administrative status is UP
```

```
Operational status is UP
```

```
Local Router ID is 10.0.0.1
```

```
Route Distinguisher is 10.0.0.1:1000 (auto-assigned)
```

```
Route Target is: 65535:268536456 (auto-assigned)
```

```
Broadcast/Unknown Unicast/Multicast traffic  
is forwarded in Ingress Replication mode
```

```
Leaf_1#show vxlan tunnels
```

```
test_vxlan
```

Destination	Source	Status
101000:10.0.0.2	101000:10.0.0.1	Up
101000:10.0.0.3	101000:10.0.0.1	Up

```
Leaf_1#show vxlan tunnels test_vxlan
```

```
test_vxlan
```

Destination	Source	Status
101000:10.0.0.2	101000:10.0.0.1	Up
101000:10.0.0.3	101000:10.0.0.1	Up

7.2.2 BGP information output

```
Leaf_1#show ip bgp l2vpn evpn
```

```
BGP table version is 9, local router ID is 10.0.0.1
```

```
Status codes: * - valid, > - best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route distinguisher: 10.0.0.1:1000					
*> [3][0][32][10.0.0.1]/88	0.0.0.0	0	100	0	?
Route distinguisher: 10.0.0.2:1000					
*>i[3][0][32][10.0.0.2]/88	10.0.0.2	0	100	0	?
Route distinguisher: 10.0.0.2:1000					
* i[3][0][32][10.0.0.2]/88	10.0.0.2	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
*>i[3][0][32][10.0.0.3]/88	10.0.0.3	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
* i[3][0][32][10.0.0.3]/88	10.0.0.3	0	100	0	?

The first entry looks as follows:

Network	Nexthop	Metric	LocPrf	Weight	Path
Route distinguisher: 10.0.0.1:1000					
*> [3][0][32][10.0.0.1]/88	0.0.0.0	0	100	0	?

Where:

- [3] – route type;
- [0] – EthTag. Not used in the current firmware version;
- [32] – length of the source route IP address;
- [10.0.0.1] – route source IP address;
- [88] – route length.

This entry indicates the creation of a local route of type 3, which is necessary for receiving BUM traffic using the ingress-replication method. This route is also advertised to all BGP neighbors with AF I2vpn evpn enabled.

The remaining entries indicate the presence of type 3 routes to remote VTEPs.

There are two instances of each route, as there are two alternative paths to each other Leaf via two Spines.

In addition, information about tunnels intended for distributing BUM traffic can be seen in the output of the **show evpn inclusive-multicast** command. Output example:

```
Leaf_1#show evpn inclusive-multicast
```

```
VXLAN test_vxlan
```

IP Address	VNI	Source	Tunnel Address	Remote Route Distinguisher	Tunnel Type
10.0.0.1	101000	Local	10.0.0.1	10.0.0.1:1000	Ingress Repl.
10.0.0.2	101000	Remote	10.0.0.2	10.0.0.2:1000	Ingress Repl.
10.0.0.3	101000	Remote	10.0.0.3	10.0.0.3:1000	Ingress Repl.

7.2.3 Checking VXLAN operation

To check the operation of the created VXLAN, check the IP connectivity between Host1 and Host2. Make sure their IP interfaces are in the same subnet.

After successfully executing the ping command, check the presence of the studied MAC addresses in the Leaf_1 and Leaf_3 tables. Example:

```
Leaf_1#show mac address-table
Flags: I - Internal usage VLAN
Aging time is 300 sec
```

Vlan	Mac Address	Interface	Type
1	e0:d9:e3:26:d6:00	0	self
1000	0c:9d:92:61:9f:c4	te1/0/11	dynamic
1000	e0:d9:e3:a8:45:40	10.0.0.3	evpn-vxlan
hu1/0/1(I)	cc:9d:a2:53:d6:80	hu1/0/1	dynamic
hu1/0/1(I)	cc:9d:a2:53:d6:81	hu1/0/1	dynamic
hu1/0/2(I)	e0:d9:e3:17:6b:40	hu1/0/2	dynamic
hu1/0/2(I)	e0:d9:e3:17:6b:41	hu1/0/2	dynamic

```
Leaf_3#show mac address-table
Flags: I - Internal usage VLAN
Aging time is 300 sec
```

Vlan	Mac Address	Interface	Type
1	e0:d9:e3:d7:ea:80	0	self
1000	0c:9d:92:61:9f:c4	10.0.0.1	evpn-vxlan
1000	e0:d9:e3:a8:45:40	te1/0/11	dynamic
hu1/0/1(I)	cc:9d:a2:53:d6:80	hu1/0/1	dynamic
hu1/0/1(I)	cc:9d:a2:53:d6:83	hu1/0/1	dynamic
hu1/0/2(I)	e0:d9:e3:17:6b:40	hu1/0/2	dynamic
hu1/0/2(I)	e0:d9:e3:17:6b:43	hu1/0/2	dynamic

The MAC address of Host1 must be present in the Leaf_3 table, type evpn-vxlan. Similarly, the MAC address of Host2 must be in the Leaf_1 table, type evpn-vxlan.

The BGP protocol output contains a type 2 route with the learned MAC address of the remote host specified as the route destination address:

```
Leaf_1#show ip bgp l2vpn evpn
```

```
BGP table version is 14, local router ID is 10.0.0.1
```

```
Status codes: * - valid, > - best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route distinguisher: 10.0.0.1:1000					
*> [2][0][0][48][0c:9d:92:61:9f:c4][0][0.0.0.0]/216	0.0.0.0	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
*>i [2][0][0][48][e0:d9:e3:a8:45:40][0][0.0.0.0]/216	10.0.0.3	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
* i [2][0][0][48][e0:d9:e3:a8:45:40][0][0.0.0.0]/216	10.0.0.3	0	100	0	?
Route distinguisher: 10.0.0.1:1000					
*> [3][0][32][10.0.0.1]/88	0.0.0.0	0	100	0	?
Route distinguisher: 10.0.0.2:1000					
*>i [3][0][32][10.0.0.2]/88	10.0.0.2	0	100	0	?
Route distinguisher: 10.0.0.2:1000					
* i [3][0][32][10.0.0.2]/88	10.0.0.2	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
*>i [3][0][32][10.0.0.3]/88	10.0.0.3	0	100	0	?
Route distinguisher: 10.0.0.3:1000					
* i [3][0][32][10.0.0.3]/88	10.0.0.3	0	100	0	?

Example of entry:

Network	Nexthop	Metric	LocPrf	Weight	Path
Route distinguisher: 10.0.0.3:1000					
*>i [2] [0] [0] [48] [e0:d9:e3:a8:45:40] [0] [0.0.0.0] /216					
	10.0.0.3	0	100	0	?

Where:

- [2] — route type;
- [0] — ESI (Ethernet segment identifier);
- [0] — EthTag. Not used in the current firmware version;
- [48] — length of the MAC address;
- [e0:d9:e3:a8:45:40] — MAC address learned on a remote VTEP;
- [0] — length of the IP address. Not used in the current firmware version;
- [0.0.0.0] — IP address. Not used in the current firmware version;
- 216 — route length.

This entry indicates the presence of a route of type 2. The e0:d9:e3:a8:45:40 MAC address is learned on a remote VTEP. For sending packets using this route, use nexthop 10.0.0.3.

Basic information on MAC addresses learned in VXLAN can also be obtained using the show **show evpn mac-ip** command:

```
Leaf_1#show evpn mac-ip
```

VXLAN test_vxlan					
VNI	VLAN	MAC Address	IP	ESI	Next Hop
101000	1000	0c:9d:92:61:9f:c4	-	-	te1/0/11
101000	1000	e0:d9:e3:a8:45:40	-	-	10.0.0.3

Traffic mirroring

Traffic mirroring can be used as a control mechanism on any device and at any time. However, in case of encapsulated traffic in VXLAN, the following information should be taken into account: outgoing traffic is intercepted by the mirror before encapsulation, so traffic encapsulated in VXLAN should only be mirrored on the device for which it is incoming.

7.3 ARP suppression

ARP suppression reduces BUM traffic on the network by reducing the number of ARP packets forwarded via the IP fabric.

Enabled by the following command in vxlan context:

```
Leaf_1#configure
Leaf_1(config)#vxlan test_vxlan
Leaf_1(config-vxlan)#arp-suppression
```

The feature is provided by a special cache that stores MAC and IP address mappings. Example of the contents of this cache:

```
Leaf_1#show arp suppression-cache
```

```
Total number of entries: 2
ARP suppression-cache timeout is 300 sec
```

IP address	Vtep	MAC address	VLAN	Port	Flags	Age
192.168.13.1	0.0.0.0	e0:d9:e3:d7:ea:80	1000	te1/0/11	local	00:00:33
192.168.13.2	10.0.0.3	e0:d9:e3:26:d6:00	1000	0	remote	00:00:08

Entries are added to the cache:

- When receiving an ARP request in a VLAN bound to VXLAN (for local type entries);
- When receiving a BGP update with a 2 MAC-IP address type (for remote type entries).

Operating principle: Host1 sends an ARP request to learn the MAC address of Host2. If MAC-IP Host2 mapping is already learned by Leaf_1 (there is a corresponding entry in a cache), then Leaf_1 will respond to the ARP request on behalf of Host2 without forwarding it through the IP fabric.

Entries are removed from the cache:

- When receiving a BGP update with a 2 MAC-IP withdraw type (for remote type entries);
- When a MAC address is removed from the device, the corresponding local entry changes to *local, inactive*. At this point, a BGP update containing the withdraw of the corresponding MAC-IP mapping is sent. The entry can become active when the MAC address is learned again. If this does not happen, the entry will be deleted upon expiration of half suppression-cache timeout period;
- Upon expiration of the suppression-cache timeout.

The suppression-cache timeout sets the maximum lifetime of local entries in the arp suppression-cache table:

```
Leaf_2(config)#arp suppression-cache timeout
<30-400000000> Seconds
```


Default value is 5 minutes.

7.4 Expected result

- VXLAN is successfully created;
- The necessary information on VXLAN and the MAC addresses learned can be found in the output of the corresponding show commands;
- User traffic is transmitted between ports of different Leaves via a VXLAN tunnel;
- ARP suppression function allows reducing the number of ARP protocol packets transmitted.

8 Configuring overlay. Multicast VXLAN

The network topology is similar to [Configuring overlay. VXLAN](#).



The initial device configurations used in this section are taken from [Configurations for OSPF](#). Additional settings required for Multicast VXLAN operation are described later in this section. Full configurations are given in Appendix 1.

8.1 Configuring Spine

Configuring PIM:

```
Spine_1(config)#interface range HundredGigabitEthernet1/0/1-3
Spine_1(config-if-range)#ip pim
Spine_1(config-if-range)#exit
Spine_1(config)#interface loopback2
Spine_1(config-if)# ip address 10.100.100.100 255.255.255.255
Spine_1(config-if)# description RP_IP
Spine_1(config-if)#exit
Spine_1(config)#ip multicast-routing pim
Spine_1(config)#ip pim rp-address 10.100.100.100
```

⚠ RP address is the same on all Spine devices.

Configuring MSDP:

```
Spine_1(config)#router msdp
Spine_1(config-msdp)#connect-source 10.0.1.1
Spine_1(config-msdp)#originator-ip 10.100.100.100
Spine_1(config-msdp)#peer 10.0.2.2
Spine_1(config-peer)#mesh-group TESTGR
Spine_1(config-peer)#exit
Spine_1(config-msdp)#exit
```

Where:

- connect-source — source address for establishing a connection with MSDP peers (loopback 1 address on Spine);
- originator-ip — IP address used as an RP address in source-active messages (loopback 2 address on Spine);
- mesh-group — name of the fully connected group. Must be the same for all Spines of the fabric.

Other Spine devices are configured in the same way, with the necessary changes made according to the diagram and IP addressing plan.

The following parameters are subject to change from device to device:

- connect-source of MSDP;
- IP address of the MSDP peer protocol.


8.2 Configuring Leaf

Configuring PIM:

```
Leaf_1(config)#interface range HundredGigabitEthernet1/0/1-2
Leaf_1(config-if-range)#ip pim
Leaf_1(config-if-range)#exit
Leaf_1(config)#interface loopback 1
Leaf_1(config-if)#ip pim
Leaf_1(config-if)#exit
Leaf_1(config)#ip multicast-routing pim
Leaf_1(config)#ip pim rp-address 10.100.100.100
Leaf_1(config)#ip multicast multipath group-paths-num
```

Other Leaf devices are configured in the same way. The commands mentioned above do not change from device to device.

The **ip multicast multipath group-paths-num** command enables PIM Join packet balancing towards available RPs. It also enables a balancing method in which the hash function calculated based on the group address is divided by N, where N is the number of available RPs.

 The above method is necessary for correct balancing when using EVPN/VXLAN. In practice, it leads to VTEP 'synchronization' and the selection of the same RP for sending traffic for a specific group.

The result of balancing will be covered in more detail in the next section.

8.3 Configuring Multicast VXLAN

VXLAN Multicast mode provides the ability to automatically search for remote VTEPs. BUM traffic replication is not performed on the outgoing VTEP, as in the default VXLAN ingress replication mode, but via PIM Multicast, closer to the destination point. Applied to the topology under consideration, this means replication not on outgoing Leaf switches, but on Spine switches.

On Leaf_1, four different VXLAN instances will be created.

Example of VXLAN instance configuration:

```
Leaf_1(config)#vlan database
Leaf_1(config-vlan)#vlan 2-5
Leaf_1(config-vlan)#exit
Leaf_1(config)#vxlan mcast2
Leaf_1(config-vxlan)# vni 102
Leaf_1(config-vxlan)# vlan 2
Leaf_1(config-vxlan)# mcast-group 233.0.0.2
Leaf_1(config-vxlan)#exit
Leaf_1(config)#vxlan mcast3
Leaf_1(config-vxlan)# vni 103
Leaf_1(config-vxlan)# vlan 3
Leaf_1(config-vxlan)# mcast-group 233.0.0.3
Leaf_1(config-vxlan)#exit
Leaf_1(config)#vxlan mcast4
Leaf_1(config-vxlan)# vni 104
Leaf_1(config-vxlan)# vlan 4
Leaf_1(config-vxlan)# mcast-group 233.0.0.4
Leaf_1(config-vxlan)#exit
Leaf_1(config)#vxlan mcast5
Leaf_1(config-vxlan)# vni 105
Leaf_1(config-vxlan)# vlan 5
Leaf_1(config-vxlan)# mcast-group 233.0.0.5
Leaf_1(config-vxlan)#exit
```

The **mcast-group <IP>** command in the VXLAN context enables BUM traffic replication in a given VXLAN via PIM Multicast.

BUM traffic arriving to VXLAN instance VLAN is intercepted by the CPU and encapsulated in a PIM Register message for registration on RP.

After registration, a VXLAN tunnel to RP is established, and BUM traffic is sent encapsulated into VXLAN header with multicast destination IP in UDP header.

- ⚠ All Leaves within the same VXLAN must use the same replication method (and the same group address in the case of multicast replication).
One mcast-group can be assigned to several VXLAN instances. 256 unique multicast groups are available.
The **mcast-group <IP>** command disables sending type 3 EVPN routes within the VXLAN.

Configuring loopback interface:

```
Leaf_1(config)#interface loopback1
Leaf_1(config-if)#ip igmp static-group 233.0.0.2
Leaf_1(config-if)#ip igmp static-group 233.0.0.3
Leaf_1(config-if)#ip igmp static-group 233.0.0.4
Leaf_1(config-if)#ip igmp static-group 233.0.0.5
```

The **ip igmp static-group** command allows Leaf to subscribe to the specified group and start listening to it. At this point, a join message is sent to this group in the direction of the RP, taking into account the balancing mechanism enabled by the **ip multicast multipath group-paths-num** command. Join messages are distributed among the RPs. On each Spine (RP), (*, G) entries are created for corresponding groups. In practice, it will look as follows:

```
Spine_1#show ip mroute
IP Multicast Routing Table

Flags: D - Dense, S - Sparse, X - IGMP Proxy, s - SSM Group,
       C - Connected, L - Local, R - RP-bit set, F - Register flag,
       T - SPT-bit set, I - Received Source Specific Host Report
Timers: Uptime/Expires

(*, 233.0.0.3), uptime: 00:00:04, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1

(*, 233.0.0.5), uptime: 00:00:03, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1
```

Leaf_1 sent to Spine_1 PIM join messages to 233.0.0.3 and 233.0.0.5 groups.

Messages to 233.0.0.2 and 233.0.0.4 groups were sent to Spine_2:

```
Spine_2#show ip mroute
IP Multicast Routing Table

Flags: D - Dense, S - Sparse, X - IGMP Proxy, s - SSM Group,
       C - Connected, L - Local, R - RP-bit set, F - Register flag,
       T - SPT-bit set, I - Received Source Specific Host Report
Timers: Uptime/Expires

(*, 233.0.0.2), uptime: 00:00:09, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1

(*, 233.0.0.4), uptime: 00:00:07, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1
```

Traffic from different groups will reach a Leaf via different Spines, thus ensuring load balancing.

Other Leaves will be configured the same way.

Check multicast route tables on both Spines:

```
Spine_1#show ip mroute
IP Multicast Routing Table

Flags: D - Dense, S - Sparse, X - IGMP Proxy, s - SSM Group,
       C - Connected, L - Local, R - RP-bit set, F - Register flag,
       T - SPT-bit set, I - Received Source Specific Host Report
Timers: Uptime/Expires

(*, 233.0.0.3), uptime: 00:00:37, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1, hu1/0/2, hu1/0/3

(*, 233.0.0.5), uptime: 00:00:37, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1, hu1/0/2, hu1/0/3
```

```
Spine_2#sh ip mroute
IP Multicast Routing Table

Flags: D - Dense, S - Sparse, X - IGMP Proxy, s - SSM Group,
       C - Connected, L - Local, R - RP-bit set, F - Register flag,
       T - SPT-bit set, I - Received Source Specific Host Report
Timers: Uptime/Expires

(*, 233.0.0.2), uptime: 00:01:39, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1, hu1/0/2, hu1/0/3

(*, 233.0.0.4), uptime: 00:01:39, expires: never, RP 10.100.100.100, Flags: S
  Incoming interface: Null, RPF neighbor 10.100.100.100
  Outgoing interface list: hu1/0/1, hu1/0/2, hu1/0/3
```

Once the other Leaves are configured, they also send PIM join messages, distributing them between the Spines (between RPs). Thus, the hu1/0/2 and hu1/0/3 interfaces on both Spines are added to the Outgoing interface list of both groups.

Now all that remains is to configure the ports on the Leaf switches for connecting client devices:

```
Leaf_1(config)#interface TenGigabitEthernet1/0/9
Leaf_1(config)#description Host1_mcast
Leaf_1(config-if)#switchport mode trunk
Leaf_1(config-if)#switchport trunk allowed vlan add 2-5
Leaf_1(config-if)#switchport forbidden default-vlan

Leaf_3(config)#interface TenGigabitEthernet1/0/9
Leaf_3(config-if)#switchport mode trunk
Leaf_3(config-if)#switchport trunk allowed vlan add 2-5
Leaf_3(config-if)#switchport forbidden default-vlan
```

8.4 Checking Multicast VXLAN configuration

To create Multicast VXLAN tunnels, send BUM traffic to client VLANs. In the example, VLAN 2–5 broadcast traffic will be sent to a client port of Leaf_1.

The result of creating these tunnels in the **show vxlan tunnels** command will be available on Leaf_2 and Leaf_3, since in this example, broadcast traffic is unidirectional and only Leaf_2 and Leaf_3 will be aware of the existence of another VTEP on the network (Leaf_1).

As a result, the output of the corresponding show commands will contain not only the result of creating VXLANs locally on VTEP, but also the result of creating tunnels.

Commands and their outputs are shown below.

- **show vxlan** — information on all VXLAN instances in a table;
- **show vxlan WORD<1-64>** — detailed information on a specific VXLAN;
- **show vxlan tunnels** — all established VXLAN tunnels;
- **show vxlan tunnels WORD<1-64>** — established VXLAN tunnels for a specific VXLAN.

Output examples:

Leaf_3#show vxlan

Name	VNI	VLAN ID	Status	BUM Forwarding	Route Distinguisher
			admin/oper		
test_vxlan	101000	1000	UP/UP	Ingress Repl.	10.0.0.3:1000
mcast2	102	2	UP/UP	Multicast VxLAN	10.0.0.3:2
mcast3	103	3	UP/UP	Multicast VxLAN	10.0.0.3:3
mcast4	104	4	UP/UP	Multicast VxLAN	10.0.0.3:4
mcast5	105	5	UP/UP	Multicast VxLAN	10.0.0.3:5

Leaf_3#show vxlan mcast2

mcast2

VxLAN Network ID is 102, VLAN ID is 2

Administrative status is UP

Operational status is UP

Local Router ID is 10.0.0.3

Route Distinguisher is 10.0.0.3:2 (auto-assigned)

Route Target is: 65500:268435558 (auto-assigned)

Broadcast/Unknown Unicast/Multicast traffic

is forwarded in multicast vxlan mode

Multicast Group address is 233.0.0.2

Leaf_3#show vxlan tunnels

test_vxlan

Destination	Source	Status
101000:10.0.0.1	101000:10.0.0.3	Up
101000:10.0.0.2	101000:10.0.0.3	Up

mcast2

Destination	Source	Status
102:10.0.0.1	102:10.0.0.3	Up

mcast3

Destination	Source	Status
103:10.0.0.1	103:10.0.0.3	Up

mcast4

Destination	Source	Status
104:10.0.0.1	104:10.0.0.3	Up

mcast5

Destination	Source	Status
105:10.0.0.1	105:10.0.0.3	Up

Leaf_3#show vxlan tunnels mcast2

mcast2

Destination	Source	Status
-----	-----	-----
102:10.0.0.1	102:10.0.0.3	Up

Broadcast traffic sent to Leaf_1 can be observed at the Leaf_3 output in unchanged form.

To make the example demonstrating the BUM traffic replication more clear, it is recommended to leave broadcasting for one VXLAN multicast group, i.e. send client broadcast traffic within one VLAN.

Port utilization on Leaf_1:

Leaf_1#show int utilization					
Port	Period, s	Sent, Kbit/s	Recv, Kbit/s	Frames sent	Frames recv
-----	-----	-----	-----	-----	-----
te1/0/9	15	0	1103	0	15029
...					
hu1/0/1	15	4	4	115	112
hu1/0/2	15	1476	4	15145	111

Outgoing traffic is greater than incoming traffic due to the addition of an extra header when encapsulating packets.

Next, traffic from this group is sent to Spine_2. Port utilization on Spine_2:

Spine_2#sh int utilization					
Port	Period, s	Sent, Kbit/s	Recv, Kbit/s	Frames sent	Frames recv
-----	-----	-----	-----	-----	-----
hu1/0/1	15	4	1476	113	15312
hu1/0/2	15	1476	4	15310	115
hu1/0/3	15	1476	4	15310	114

As can be seen, traffic is copied to other Leaves on Spine_2.

Traffic mirroring


Traffic mirroring can be used as a control mechanism on any device and at any time. However, in case of encapsulated traffic in VXLAN, the following information should be taken into account: outgoing traffic is intercepted by the mirror before encapsulation, so traffic encapsulated in VXLAN should only be mirrored on the device for which it is incoming.

8.5 Expected result

- PIM Multicast mode is successfully used;
- Information about the operating mode and tunnels established using PIM multicast can be viewed in the output of the corresponding show commands;
- BUM traffic replication using PIM multicast is implemented on Spine switches.

9 Configuring overlay. Symmetric IRB

The network topology is shown below.



The initial device configurations used in this section (without VXLANs created) are taken from [Configurations for OSPF](#).

Additional settings required for L3VNI operation are described later in this section.

Full configurations are given in Appendix 1.

9.1 Configuring Symmetric IRB

For EVPN/VXLAN, L3VPN configuration using Symmetric IRB is supported, specifically routing via L3VNI.

Create L3VNIs on all Leaves and prepare them for operation.

Example of the L3VNI creation and configuration commands execution on Leaf_1:

```
Leaf_1(config)#vlan database
Leaf_1(config-vlan)#vlan 100
Leaf_1(config-vlan)#exit
Leaf_1(config)#ip vrf VRF1
Leaf_1(config-vrf)#vni 100100
Leaf_1(config-vrf)#route-target both 65500:100100
Leaf_1(config-vrf)#exit
Leaf_1(config)#vxlan L3_vxlan
Leaf_1(config-vxlan)#vni 100100 ip-routing
Leaf_1(config-vxlan)#vlan 100
Leaf_1(config-vxlan)#exit
Leaf_1(config)#interface vlan 100
Leaf_1(config-if)#ip vrf VRF1
Leaf_1(config-if)#exit
```


Where:

- vni 100100 – L3VNI. Must be created on all VTEPs where routing is required. Only one L3VNI can be assigned to a single VRF. L3VNI values must not be repeated in different VRFs;
- route-target both 65500:100100 – RT community value. This extended community is added to outgoing update messages with type 2 (ip-mac) and 5 routes. When receiving such routes, the RT community value determines which VRF the received routes are set to. RT values must not be repeated in different VRFs;
- vxlan L3_vxlan – VXLAN required to enable routing in L3VNI;
- vni 100100 ip-routing – parameter indicating that vni 100100 will be used for routing;
- vlan 100 – must be used in conjunction with L3VNI. This VLAN will not be used to transmit any traffic, i.e., it does not need to be assigned to any interface.
It is necessary to assign the interface of this VLAN to the corresponding VRF.

Additional configuration for the BGP protocol:

```
Leaf_1(config)#router bgp
Leaf_1(router-bgp)#vrf VRF1
Leaf_1(router-bgp-vrf)#address-family ipv4 unicast
Leaf_1(router-bgp-af-vrf)#redistribute connected
Leaf_1(router-bgp-af-vrf)#exit
Leaf_1(router-bgp-vrf)#exit
Leaf_1(router-bgp)#exit
```

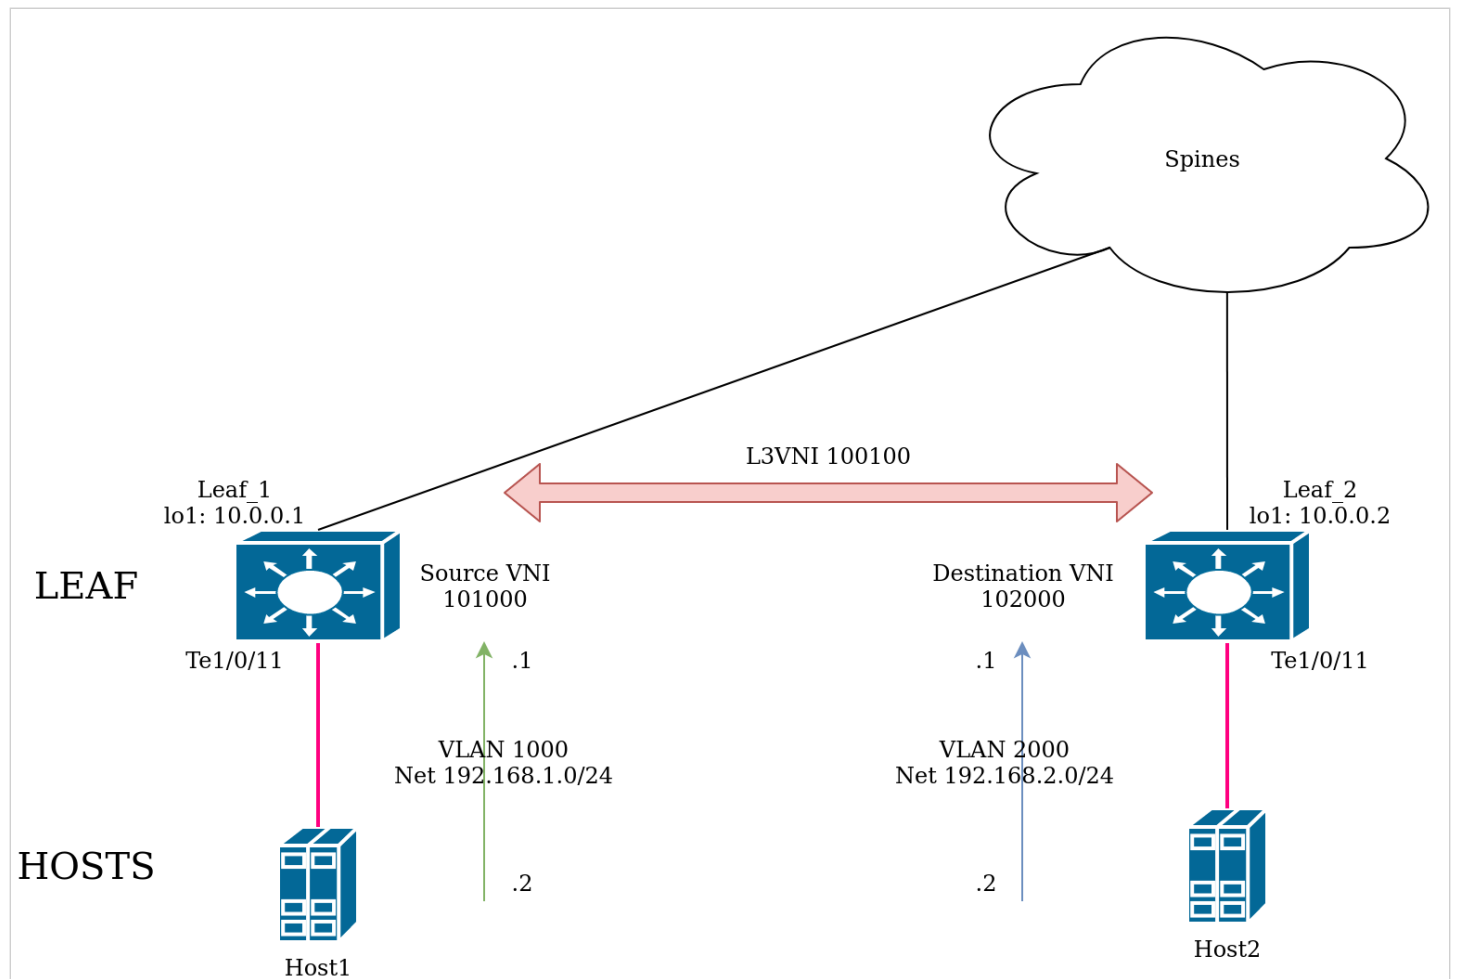
This setting launches redistribution of routes to connected networks to the created L3VPN. In this example, organized in VRF1.

 It is also possible to redistribute routes from the OSPF protocol and from eBGP neighbors in L3VPN. To do this, create and launch an OSPF protocol instance in the corresponding VRF, which will operate at the junction of the IP fabric and the network connected to it. In case of eBGP, establish a session with a BGP neighbor in the corresponding VRF. Examples of Leaf configuration for OSPF and eBGP are given in the symmetric IRB section of Appendix 1.

Configured redistribution is required to send type 5 routes announcing the availability of client networks in a given VRF.

Other Leaf devices are configured in the same way.

The configured routing scheme is shown below.



On Leaf_1, a VXLAN with VNI 101000 and VLAN 1000 are created. The VLAN 1000 IP interface is added to VRF1 so that L3VNI is used to route client packets. Then an IP address is assigned to this IP interface.

```
Leaf_1(config)#vlan database
Leaf_1(config-vlan)#vlan 1000
Leaf_1(config-vlan)#exit
Leaf_1(config)#vxlan test_vxlan1
Leaf_1(config-vxlan)#vni 101000
Leaf_1(config-vxlan)#vlan 1000
Leaf_1(config-vxlan)#exit
Leaf_1(config)#interface vlan 1000
Leaf_1(config-if)#ip vrf VRF1
Leaf_1(config-if)#ip address 192.168.1.1 /24
Leaf_1(config-if)#exit
Leaf_1(config)#interface TenGigabitEthernet1/0/11
Leaf_1(config-if)#description Host1
Leaf_1(config-if)#switchport access vlan 1000
Leaf_1(config-if)#exit
```

On Leaf_2, a VXLAN with VNI 102000 and VLAN 2000 are created. The VLAN 2000 IP interface is added to VRF1.

```

Leaf_2(config)#vlan database
Leaf_2(config-vlan)#vlan 2000
Leaf_2(config-vlan)#exit
Leaf_2(config)#vxlan test_vxlan2
Leaf_2(config-vxlan)#vni 102000
Leaf_2(config-vxlan)#vlan 2000
Leaf_2(config-vxlan)#exit
Leaf_2(config)#interface vlan 2000
Leaf_2(config-if)#ip vrf VRF1
Leaf_2(config-if)#ip address 192.168.2.1 /24
Leaf_2(config-if)#exit
Leaf_2(config)#interface TenGigabitEthernet1/0/11
Leaf_2(config-if)#description Host2
Leaf_2(config-if)#switchport access vlan 2000
Leaf_2(config-if)#exit

```

9.2 Checking Symmetric IRB configuration

To verify the functionality of the scheme, check the L3 connectivity between hosts. For example, execute the ping command from Host1 to Host2, provided that the hosts have routes to each other's networks.

Traffic sent between client networks can be seen using mirroring on Spine_1 (or on both Spines in case of heterogeneous traffic balancing) encapsulated in VXLAN using configured L3VNI 100100.

Traffic mirroring

Traffic mirroring can be used as a control mechanism on any device and at any time. However, in case of encapsulated traffic in VXLAN, the following information should be taken into account: outgoing traffic is intercepted by the mirror before encapsulation, so traffic encapsulated in VXLAN should only be mirrored on the device for which it is incoming.

To view route information in a specific VRF, use the **show ip route vrf WORD<1-32>** command.


Example of the command execution:

```

Leaf_1#show ip route vrf VRF1
Maximum Parallel Paths: 32 (32 after reset)
Load balancing: src-dst-mac-ip
IP Forwarding: enabled
Codes: > - best, C - connected, S - static,
        R - RIP,
        O - OSPF intra-area, OIA - OSPF inter-area,
        OE1 - OSPF external 1, OE2 - OSPF external 2,
        B - BGP, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, ia - IS-IS inter area
[d/m]: d - route's distance, m - route's metric

C    192.168.1.0/24 is directly connected, vlan 1000
B    192.168.2.0/24 [200/0] via 10.0.0.2, 00:04:26, VNI 100100,
router-mac e0:d9:e3:17:6b:40
B    192.168.2.2/32 [200/0] via 10.0.0.2, 00:04:26, VNI 100100,
router-mac e0:d9:e3:17:6b:40

```

 The route to the 192.168.2.2/32 network, i.e. to Host2, was obtained by a type 2 IP-MAC announcement. This announcement was sent from Leaf_2 while learning about Host2 via ARP. This route is necessary for routing traffic to a specific host if the client network is located behind several Leaves.

The route to the 192.168.2.0/24 network specifies router-mac. This is the MAC address of the destination router. In the case under consideration, it is Leaf_2.

This route is obtained via a type 5 announcement.

Type 5 routes can be seen in the BGP protocol information output. The example below shows a shortened version of the output containing only type 5 routes:

```
Leaf_1#show ip bgp l2vpn evpn

BGP routing table information for VRF default
BGP table version is 10, local router ID is 10.0.0.1
Status codes: * - valid, > - best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network                Nexthop          Metric    LocPrf    Weight    Path
   .
   .
   .
Route distinguisher: 10.0.0.1:34564
*> [5][0][0][24][192.168.1.0]/224
           0.0.0.0           0           100         0           ?

Route distinguisher: 10.0.0.2:34564
*>i[5][0][0][24][192.168.2.0]/224
           10.0.0.2           0           100         0           ?

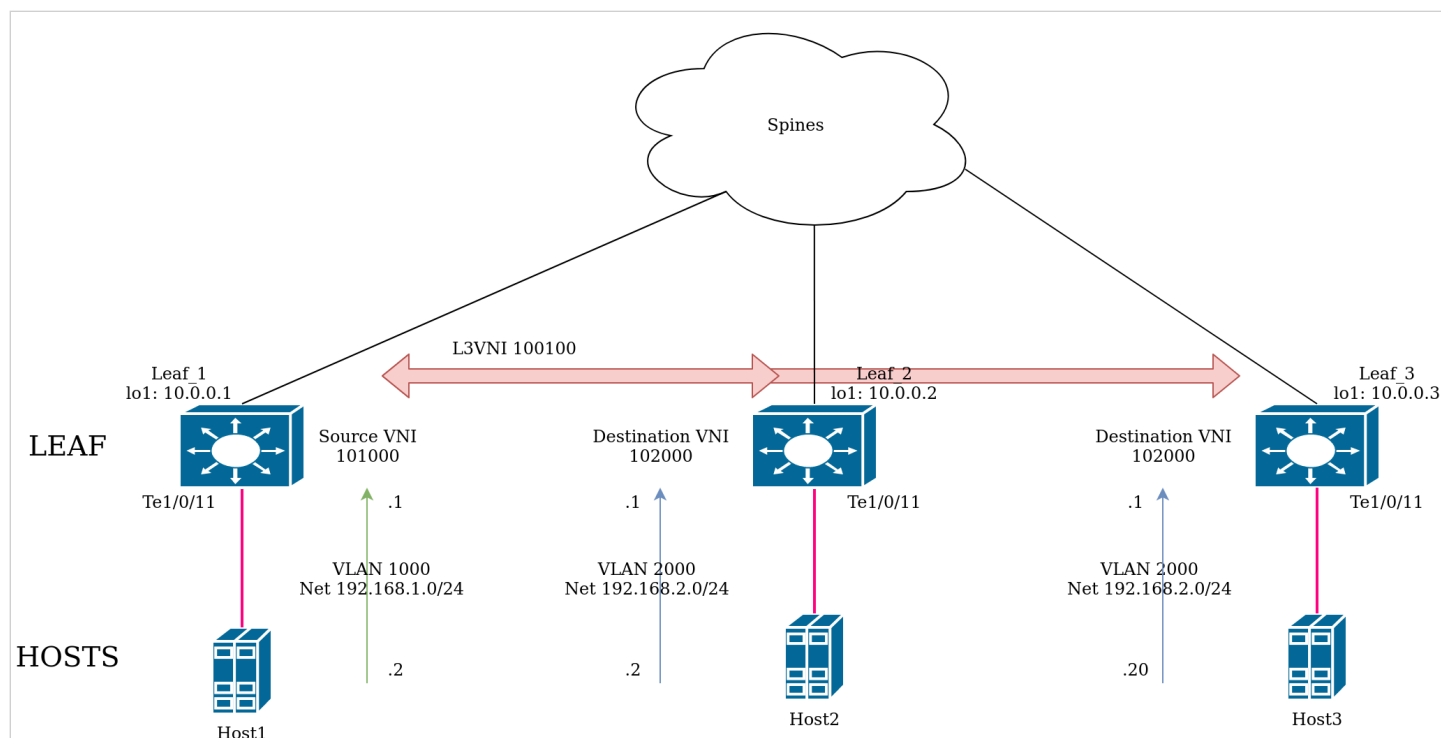
Route distinguisher: 10.0.0.2:34564
* i[5][0][0][24][192.168.2.0]/224
           10.0.0.2           0           100         0           ?
```

Where:

- [5] — route type;
- [0] — ESI (Ethernet segment identifier);
- [0] — EthTag. Not used in the current firmware version;
- [24] — prefix length;
- [192.168.2.0] — destination network IP address;
- 224 — route length.

9.3 Anycast gateway

Complete the logical diagram as shown below.



Anycast-gateway complements the ability of different Leaves to act as gateways for the same network. When using anycast-gateway, hosts do not need to update their ARP entries when switching from one Leaf to another (for example, when migrating virtual machines), because gateway MAC address does not change from Leaf to Leaf due to the configuration. It is suitable for both multihomed connection networks and cases where the same network is divided into several physical segments connected to different Leaves.

Additional configuration for Leaf_3. VXLAN on Leaf_3 is completely identical to Leaf_2, so the same configuration and the same network (192.168.2.0/24) can be used.

```
Leaf_3(config)#vlan database
Leaf_3(config-vlan)#vlan 2000
Leaf_3(config-vlan)#exit
Leaf_3(config)#vxlan test_vxlan2
Leaf_3(config-vxlan)#vni 102000
Leaf_3(config-vxlan)#vlan 2000
Leaf_3(config-vxlan)#exit
Leaf_3(config)#interface vlan 2000
Leaf_3(config-if)#ip vrf VRF1
Leaf_3(config-if)#ip address 192.168.2.1 /24
Leaf_3(config-if)#exit
Leaf_3(config)#interface TenGigabitEthernet1/0/11
Leaf_3(config-if)#description Host3
Leaf_3(config-if)#switchport access vlan 2000
Leaf_3(config-if)#exit
```

Anycast-gateway is configured in to steps.

1. Set a virtual MAC address that will replace the base MAC address of the switch in ARP psckets sent from interfaces on which this feature is enabled. Example for Leaf_2 and VLAN 2000:

```
Leaf_2(config)#anycast-gateway mac-address 00:00:00:11:11:11
```

2. Enable anycast-gateway on VLAN interface:

```
Leaf_2(config)#interface vlan 2000
Leaf_2(config-if)#anycast-gateway
```

3. Similarly, configure Leaf_3 using the same virtual MAC address.

To view the settings made, use the show command:

```
Leaf_2#show ip anycast-gateway
Anycast-gateway virtual MAC address: 00:00:00:11:11:11

Anycast-gateway is configured on interfaces:

Vlans:
  2000
```

Checking anycast-gateway configuration

1. Ping all hosts. As a result:
L3 connectivity between Host1 and the other hosts;
L2 connectivity between Host2 and Host3;
2. In the ARP table of Hosts 2 and 3, there is a match between the configured gateway address and the virtual MAC address of the anycast gateway. Example of viewing the table on Host3:

```
Host3#sh arp

Total number of entries: 1
```

VLAN	Interface	IP address	HW address	status	IP Unnumbered I/F
vlan 1	te1/0/1	192.168.2.1	00:00:00:11:11:11	dynamic	

9.4 Gateway-ip for type 5 routes

This function allows setting the value of IP-v4 Gateway address field in NLRI of sent EVPN routes of type 5. Setting is performed by linking to the neighbor route-map with the corresponding set action.

⚠ Route-map with this setting should be used in the context of address-family l2vpn evpn of the BGP neighbor and only in the out direction.

Since in case of an IP fabric, Leaf devices are neighbors with Spine switches, this route-map must be configured towards all Spine switches for the IP fabric to operate correctly.

Example of Leaf_1 configuration:

```
Leaf_1(config)#route-map TESTMAP 10 permit
Leaf_1(config-route-map)#set evpn gateway-ip 192.168.50.1
Leaf_1(config-route-map)#exit
Leaf_1(config)#router bgp 65500
Leaf_1(router-bgp)#neighbor 10.0.1.1
Leaf_1(router-bgp-nbr)#address-family l2vpn evpn
Leaf_1(router-bgp-nbr-af)#route-map TESTMAP out
Leaf_1(router-bgp-nbr-af)#exit
Leaf_1(router-bgp-nbr)#exit
Leaf_1(router-bgp)#neighbor 10.0.2.2
Leaf_1(router-bgp-nbr)#address-family l2vpn evpn
Leaf_1(router-bgp-nbr-af)#route-map TESTMAP out
Leaf_1(router-bgp-nbr-af)#exit
Leaf_1(router-bgp-nbr)#exit
```


The result can be seen in the IPv4 Gateway address in NLRI of sent type 5 routes:

```
EVPN NLRI: IP Prefix route
Route Type: IP Prefix route (5)
Length: 34
Route Distinguisher: 00010a0000018704 (10.0.0.1:34564)
ESI: 00:00:00:00:00:00:00:00:00
Ethernet Tag ID: 0
IP prefix length: 24
IPv4 address: 192.168.1.0
IPv4 Gateway address: 192.168.50.1
```

Remote Leaves will add the received routes in routing tables, taking into account the specified IP-v4 Gateway address value and the availability of the specified gateway.

9.5 DHCP relay

The network topology is shown below.



The initial device configurations used in this section are taken from [Configurations for Symmetric IRB](#). Additional settings required for DHCP relay agent operation are described later in this section.

9.5.1 Configuring DHCP relay

No additional settings are required on Leaf_1 from the DHCP server side. The server will be placed in the existing VLAN 1000, TenGigabitEthernet1/0/11 interface.

⚠ Case 1. Host2 and Host3 are on the same network but on different Leaves. The DHCP server is on a different network. Accordingly, it is necessary to solve the problem of sending a response from the DHCP server to the desired Leaf, while issuing an IP address to the host from the desired address pool.

In the example below, Leaf_2 will be configured.

In VRF1, create an additional loopback interface, which will be the source interface for DHCP requests forwarded to the DHCP server. The addresses of these loopbacks must also be unique for each Leaf. This will allow responses from the DHCP server to return to the correct Leaf.

```
Leaf_2(config)#interface loopback2
Leaf_2(config-if)#ip vrf VRF1
Leaf_2(config-if)#ip address 10.0.0.102 255.255.255.255
Leaf_2(config-if)#exit
```

The DHCP server must have routes to these loopback interfaces or a default route in order to send responses to DHCP requests.

Enable the DHCP relay agent function on the VLAN that Host2 is connected to. In this case, this is VLAN 2000:

```
Leaf_2(config)#interface vlan 2000
Leaf_2(config-if)#ip dhcp relay enable
Leaf_2(config-if)#exit
```

Perform all necessary settings in global configuration mode:

```
Leaf_2(config)#ip dhcp relay address 192.168.11.2 vrf VRF1
Leaf_2(config)#ip dhcp relay enable
Leaf_2(config)#ip dhcp relay source-interface loopback 2 vrf VRF1
Leaf_2(config)#ip dhcp information option
Leaf_2(config)#ip dhcp information option vpn link-selection
```

- `ip dhcp relay address 192.168.11.2 vrf VRF1` – DHCP server address to which DHCP requests within the VRF will be redirected. When transmitted between Leaves, these packets will be encapsulated in VXLAN using L3 VNI. In this case, VNI 100100;
- `ip dhcp relay enable` – enable DHCP Relay agent function globally on the device;
- `ip dhcp relay source-interface loopback 2 vrf VRF1` – in this case, src ip and giaddr (Relay agent IP address) will be replaced with the loopback2 interface address in DHCP packets sent to the server;
- `ip dhcp information option` – enable adding option 82 to the DHCP request;
- `ip dhcp information option vpn link-selection` – enable adding suboption 5 to option 82 of the DHCP request. If only the *vpn* parameter is specified, all the suboptions (5, 11, 151) will be added.

Suboption 5

Link Selection. Contains the client's network address. Allows separating giaddr (Relay agent IP address) and the client's network, thereby enabling the DHCP server to issue IP addresses based this suboption.


Suboption 11

Server ID Override. Contains the IP address of the interface on which Relay agent is enabled. When this suboption is enabled, the device overwrites option 54 (DHCP Server Identifier) in responses from the DHCP server, with the value from suboption 11.

Suboption 151

Virtual Subnet Selection. Contains VRF name on the client interface. Can also be used for the DHCP server to assign addresses from a specific pool.

Leaf_3 is configured the same way, except for the IP address on the loopback2 interface.


 An example of configuring `isc-dhcp-server` to process suboption 5 is provided at the end of this section.

Expected result

- DHCP requests from Host2 and Host3 are intercepted by DHCP Relay agent on Leaf;
- When sent to the DHCP server, option 82 with suboption 5 will be added to the packets:

```
Option: (82) Agent Information Option
  Length: 38
  Option 82 Suboption: (1) Agent Circuit ID
  Option 82 Suboption: (2) Agent Remote ID
  Option 82 Suboption: (5) Link selection (192.168.2.0)
    Length: 4
    Link selection: 192.168.2.0
```

- The IP address of the loopback2 interface will be inserted into the Relay agent IP address field;
- DHCP server can issue IP addresses from different pools using suboption 5 of option 82;
- Responses from the DHCP server are sent to the Leaf from which the request originated, due to the unique Loopback 2 interface addresses used as the Relay agent IP address in DHCP requests.

 **Case 2.** Host2 and Host3 are on different networks. Each client network is connected to only one Leaf switch. IP addresses of all Relay agents are unique.

The interfaces towards the hosts on Leaf look as follows:

```
Leaf_2:

interface vlan 2000
 ip vrf VRF1
 ip address 192.168.2.1 255.255.255.0
 ip dhcp relay enable
 anycast-gateway
!

Leaf_3:


interface vlan 2000
 ip vrf VRF1
 ip address 192.168.3.1 255.255.255.0
 ip dhcp relay enable
 anycast-gateway
!
```

In this case, Leaf configuration can be simplified. For DHCP Relay to operate correctly, perform the following settings in global configuration mode:

```
Leaf_2(config)#ip dhcp relay address 192.168.11.2 vrf VRF1
Leaf_2(config)#ip dhcp relay enable
```

Creating additional Loopback interfaces on Leaf is not required.

The DHCP server must have routes to client networks or a default route in order to send responses to DHCP requests.

 An example of configuring isc-dhcp-server to process suboption 5 is provided at the end of this section.


Expected result

- DHCP requests from Host2 and Host3 will be intercepted on Leaf by DHCP Relay agent;
- In DHCP packets sent to the server, src ip and giaddr (Relay agent IP address) with the address configured on client interfaces from which DHCP request was sent. 192.168.2.1 for Leaf_2, and 192.168.3.1 for Leaf_3;
- DHCP server can issue IP addresses from different pools according to the Relay agent IP address field of DHCP requests.

⚠ Case 3. Host2 and Host3 are on the same network but on different Leaves. Host4 is on different network.

Topology

The network topology is shown below.




A specific VXLAN is created for Host4. The interface for the Host4 network must also belong to VRF1:

```
vxlan test_vxlan4
vni 102004
vlan 2004
exit
!
interface vlan 2004
ip vrf VRF1
ip address 192.168.4.1 255.255.255.0
ip dhcp relay enable
exit
!
```

The solution to this problem is to perform the configuration described in Case 1.

Due to the unique addresses of Loopback 2 interfaces, responses from the DHCP server will be sent to the Leaf from the Relay agent of which the request originated. Due to suboption 5 of option 82, it will be possible to issue addresses from a specific pool.

Example of isc-dhcp-server configuration for suboption 5 processing

 This example is provided for cases where the isc-dhcp-server version does not support automatic recognition of suboption 5 of option 82. Otherwise, the 'class' and 'allow members' settings are not necessary.

```
class "192.168.2.0" {
    match if binary-to-ascii (10, 8, ".", option agent.link-selection) = "192.168.2.0";
}

class "192.168.3.0" {
    match if binary-to-ascii (10, 8, ".", option agent.link-selection) = "192.168.3.0";
}

class "192.168.4.0" {
    match if binary-to-ascii (10, 8, ".", option agent.link-selection) = "192.168.4.0";
}

#-----For Leaf_2-----
subnet 192.168.2.0 netmask 255.255.255.0 {
    pool {
        allow members of "192.168.2.0";
        range 192.168.2.10 192.168.2.254;
        option subnet-mask 255.255.255.0;
        option routers 192.168.2.1;
    }
}

#-----For Leaf_3-----
subnet 192.168.3.0 netmask 255.255.255.0 {
    pool {
        allow members of "192.168.3.0";
        range 192.168.3.10 192.168.3.254;
        option subnet-mask 255.255.255.0;
        option routers 192.168.3.1;
    }
}

#-----For Leaf_3_2-----
subnet 192.168.4.0 netmask 255.255.255.0 {
    pool {
        allow members of "192.168.4.0";
        range 192.168.4.10 192.168.4.254;
        option subnet-mask 255.255.255.0;
        option routers 192.168.4.1;
    }
}
```

Classes allow selecting by the value of suboption 5 of option 82.


The condition for issuing an address from a specific pool is that the request belongs to a specific class.

9.6 Expected result

- L3VNI provides symmetric routing between client networks located on different Leaves and in different VXLANs;
- The anycast-gateway feature allows replacing MAC address of the gateway created on Leaf devices;
- It is possible to set the value of the IP-v4 Gateway address field in type 5 EVPN routes;
- DHCP Relay agent allows forwarding DHCP packets to the server and back, adding the necessary options to the request.

10 EVPN multihoming

The network topology is shown below.



The initial device configurations used in this section are taken from [Configurations for multicast VXLAN](#).

- ⚠ EVPN multihoming is considered in this section using multicast VXLAN as an example, since this option is more flexible in terms of settings. The ingress replication mode can also be used. In this case, no multicast settings are required.


Additional settings required for EVPN multihoming operation are described later in this section.

- ⚠ In order to reduce topology convergence time in case of multihomed connection to IP fabric, use reduced intervals between sending BGP update messages on all Leaf switches. They are configured in the context of the BGP neighbor configuration.

Example of BGP configuration for Leaf_1:

```
Leaf_1(config)#router bgp 65500
Leaf_1(router-bgp)#neighbor 10.0.1.1
Leaf_1(router-bgp-nbr)#advertisement-interval 1 withdraw 1
Leaf_1(router-bgp-nbr)#exit
Leaf_1(router-bgp)#neighbor 10.0.2.2
Leaf_1(router-bgp-nbr)#advertisement-interval 1 withdraw 1
```

Full configurations are given in Appendix 1.

 EVPN/VXLAN in conjunction with MLAG (VPC) is not supported in the current firmware version.

10.1 Configuring EVPN multihoming

Host1 will act as a multihomed client. Connection to Leaf_1-2 is established via port-channel using LACP.

Additional configuration for Leaf_1 and Leaf_2:

```
Leaf_1(config)#interface port-channel 1
Leaf_1(config-if)#description Host1_LAG
Leaf_1(config-if)#switchport mode trunk
Leaf_1(config-if)#switchport trunk allowed vlan add 2-5
Leaf_1(config-if)#ethernet-segment 1000
Leaf_1(config-es)#system-mac 11:22:33:44:55:66
Leaf_1(config-es)#exit
Leaf_1(config-if)#interface TenGigabitEthernet1/0/12
Leaf_1(config-if)#channel-group 1 mode auto
Leaf_1(config-if)#exit
```

Where:

- ethernet-segment 1000 — create an Ethernet segment with the number 1000;
- system-mac 11:22:33:44:55:66 — set a MAC address used as System ID of the LACP protocol.

Leaf_2 is configured the same way.

The following functions are configured by default and do not require additional settings.

- Split horizon BUM traffic filtering — prevents BUM traffic that was sent from an Ethernet segment from returning to that same segment via other Leaves connected to that segment;
- Designated forwarder — only the device selected by DF can send BUM traffic to the Ethernet segment;
- Local bias — if the source Ethernet segment and the destination Ethernet segment are connected to the same Leaf, traffic between them is not sent to the Spine;
- Fast convergence — VTEP sends a type 1 BGP withdraw message when a link is lost in the Ethernet segment. The fast convergence mechanism allows remote VTEPs to delete MAC addresses learned on this VTEP/ES when receiving a type 1 BGP withdraw.

10.2 Checking EVPN multihoming configuration

⚠ Since the Po1 interfaces on the Leaf switches are configured in trunk mode, the Host1 and Host2 interfaces must also be configured in trunk mode with VLAN 2 specified.

On Host1, check that both Leaf links were added to port-channel.

On Leaf, check the status of the configured port-channel:

```
Leaf_1#show interfaces Port-Channel1
Port-Channel1 is up (connected)
  Interface index is 1000
  Hardware is aggregated ethernet interface(s), MAC address is cc:9d:a2:53:d6:81
  Description: Host1_LAG
  Interface MTU is 9000
  Link is up for 0 days, 0 hours, 20 minutes and 13 seconds
  Link aggregation type is LACP
  No. of members in this port-channel: 1 (active 1)
    TenGigabitEthernet1/0/12, full-duplex, LACP active, 10000Mbps (active)
  Active bandwidth is 10000Mbps
  15 second input rate is 0 Kbit/s
  15 second output rate is 0 Kbit/s
    83 packets input, 7936 bytes received
    0 broadcasts, 83 multicasts
    0 input errors, 0 FCS
    0 oversize, 0 internal MAC
    0 pause frames received
    184 packets output, 45824 bytes sent
    101 broadcasts, 83 multicasts
    0 output errors, 0 collisions
    N/S excessive collisions, 0 late collisions
    0 pause frames transmitted
```

Then check the status of ethernet-segment:

```
Leaf_1#show evpn ethernet-segment
```

Interface	Status admin/oper	ES Number	Ethernet Segment ID (ESI)	Remote members
Po1	UP/UP	1000	03:11:22:33:44:55:66:00:03:e8	10.0.0.2

```
Leaf_1#show evpn ethernet-segment 03:11:22:33:44:55:66:00:03:e8 detailed
```

Ethernet Segment: Po1

ESI: 03:11:22:33:44:55:66:00:03:e8

ES number is 1000

ES system MAC address is 11:22:33:44:55:66

Administrative status is up

Operational status is up

All-Active multi-homing mode

Route Distinguisher is 10.0.0.1:0 (auto-assigned)

Route Target is 1122.3344.5566 (auto-assigned)

Designated Forwarder election delay is 3 seconds

Members:

Router ID	Type
10.0.0.1	local
10.0.0.2	remote

Connected VXLANs:

VNI	VLAN ID	Designated Forwarder	Name
102	2	10.0.0.1	mcast2
103	3	10.0.0.2	mcast3
104	4	10.0.0.1	mcast4
105	5	10.0.0.2	mcast5

The detailed information output shows the configured ES number and ES system MAC address, as well as ESI generated based on them: 03:11:22:33:44:55:66:00:03:e8.

In addition, information about all member devices of the Ethernet segment is shown. In this case, Router ID 10.0.0.1 and 10.0.0.2.

The selection of the Designated Forwarder may differ for different VXLANs, as can be seen in the Connected VXLANs table. DF is responsible for sending BUM traffic to a specific Ethernet segment within a specific VLAN.

The MAC address of Host1 must be present in the Leaf_3 MAC address table (information output is shortened for clarity of the example):

```
Leaf_3#show mac address-table
Flags: I - Internal usage VLAN
Aging time is 300 sec
```

Vlan	Mac Address	Interface	Type
2	e0:d9:e3:a8:45:40	10.0.0.2	evpn-vxlan
		10.0.0.1	

Where:

- e0:d9:e3:a8:45:40 – Host1 MAC address. The interface field contains Leaf_1 and Leaf_2 IP addresses, which indicates that this MAC address is accessible via two tunnels.

The BGP route information output will contain type 2 routes with specified ESI (the output is shortened for clarity of the example):

```
Leaf_3#show ip bgp l2vpn evpn
```

```
BGP routing table information for VRF default
BGP table version is 60, local router ID is 10.0.0.3
Status codes: * - valid, > - best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Nextthop	Metric	LocPrf	Weight	Path
Route distinguisher: 10.0.0.2:2					
*>i [2] [03:11:22:33:44:55:66:00:03:e8] [0] [48] [e0:d9:e3:a8:45:40] [0] [0.0.0.0] /216	10.0.0.2	0	100	0	?
Route distinguisher: 10.0.0.2:2					
* i [2] [03:11:22:33:44:55:66:00:03:e8] [0] [48] [e0:d9:e3:a8:45:40] [0] [0.0.0.0] /216	10.0.0.2	0	100	0	?

Where:

- [2] — route type;
- [03:11:22:33:44:55:66:00:03:e8] — ESI (Ethernet segment identifier). Created based on the configured Ethernet segment number and System MAC;
- [0] — EthTag. Not used in the current firmware version;
- [48] — MAC address length;
- [e0:d9:e3:a8:45:40] — Host1 MAC address learned on a remote VTEP;
- [0] — IP address length. Not used in the current firmware version;
- [0.0.0.0] — IP address length. Not used in the current firmware version;
- 216 — route length.

To check the operation of the multihoming connection, check IP connectivity between Host1 and Host2. Make sure their IP interfaces in VLAN2 are in the same subnet.

Disconnecting links between Host1 and Leaves one by one will not result in loss of IP connectivity.

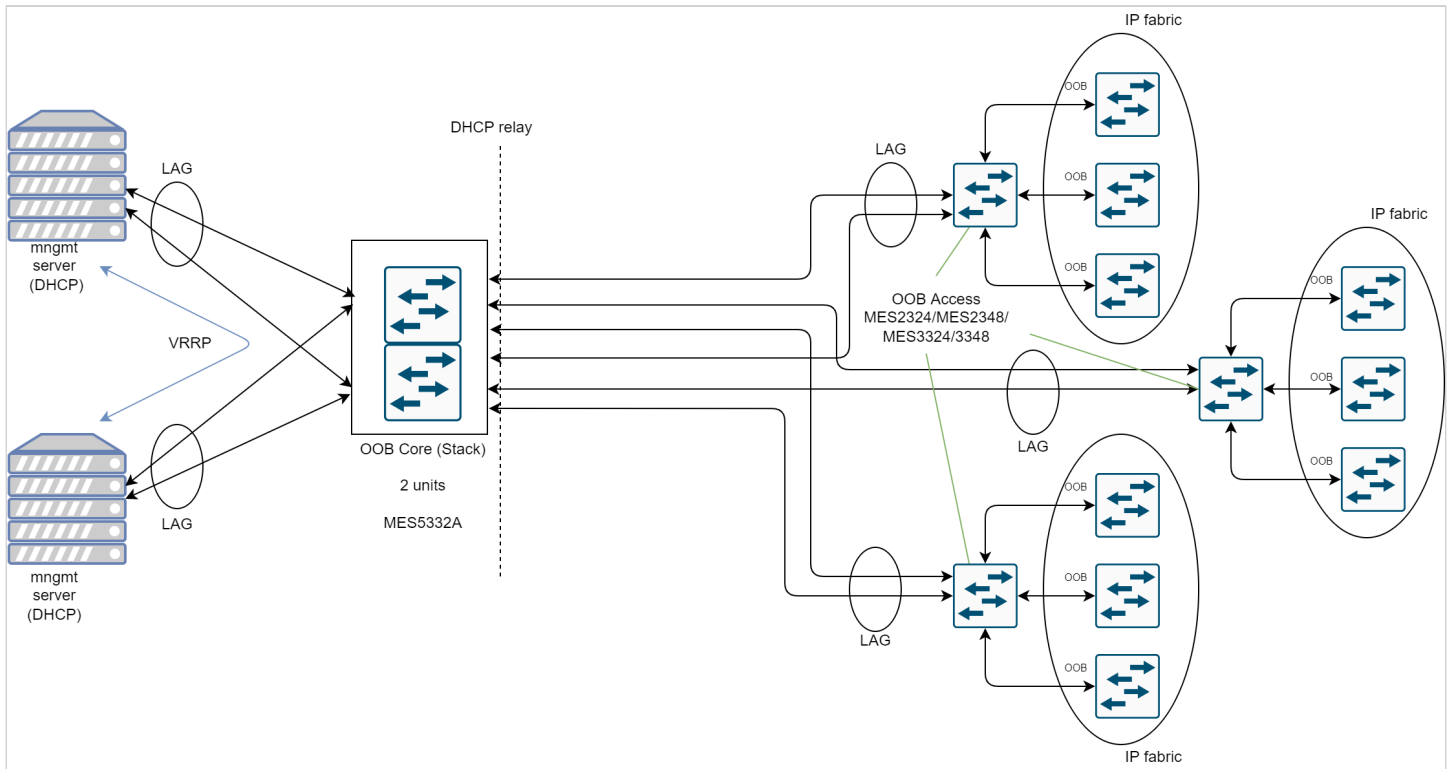
10.3 Expected result

- It is possible to connect a client device to several Leaves via LAG;
- User traffic transfer between multihomed and singlhomed clients is successful;
- Due to various BUM traffic filtering mechanisms, no loops occur when exchanging data with a multihomed client.

11 Management network

The out-of-band management approach involves transferring management information separately from data transfer. Applied to an IP fabric, this means deploying a network that provides access to the OOB interfaces of all fabric devices without intersecting with the IP fabric network itself. This approach allows devices to be managed regardless of the status of the network transmitting commercial information.

11.1 OOB network



⚠ The device models shown in the diagram are recommended.

- OOB core — a stack of aggregation switches that performs routing and DHCP relay functions. It is used to connect to the OOB access switch management network via LAG to improve fault tolerance
- OOB access switches act as switches providing access to the management network for managed (IP fabric) devices via OOB interfaces.

The diagram uses the following protocols and technologies:

- LAG (Link Aggregation group) — each port group must consist of Ethernet interfaces with the same speed operating in duplex mode. Aggregating ports into a group increases the bandwidth between interacting devices and improves fault tolerance. A port group is a single logical port for the switch.
- LACP (Link Aggregation Control Protocol) — allows aggregating several physical channels into one, thus creating the above-described LAG.
- VRRP (Virtual Router Redundancy Protocol) — allows reserving a default gateway used by all IP fabric devices to exchange information with the management, monitoring, time synchronization systems.
- DHCP Relay agent. Transfers DHCP packets from the client to the server and back if the DHCP server is on one network and the client is on another. Another function is to add additional options to client DHCP requests (for example, option 82).
- Stack — OOB core switches are stacked. Stacking allows them to function as a single device, thereby increasing fault tolerance and simplifying management. It is recommended to use a ring topology to increase the fault tolerance of the stack.

11.2 Device configurations

11.2.1 OOB core

```
no spanning-tree
!
vlan database
vlan 2,111
exit
!
ip dhcp relay address 192.168.11.1
ip dhcp relay enable
ip dhcp snooping
ip dhcp snooping vlan 2
!
hostname OOB_Core
!
interface TenGigabitEthernet1/0/1
channel-group 1 mode auto
exit
!
interface TenGigabitEthernet1/0/2
channel-group 2 mode auto
exit
!
interface TenGigabitEthernet1/0/3
channel-group 3 mode auto
exit
!
interface TenGigabitEthernet1/0/11
description Server
switchport access vlan 111
exit
!
interface TenGigabitEthernet2/0/1
channel-group 1 mode auto
exit
!
interface TenGigabitEthernet2/0/2
channel-group 2 mode auto
exit
!
interface TenGigabitEthernet2/0/3
channel-group 3 mode auto
exit
!
interface TenGigabitEthernet2/0/11
description Server
switchport access vlan 111
exit
!
interface Port-Channel1
description OOB_access_1
switchport mode trunk
switchport trunk allowed vlan add 2
switchport forbidden default-vlan
exit
!
```

```
interface Port-Channel2
description 00B_access_2
switchport mode trunk
switchport trunk allowed vlan add 2
switchport forbidden default-vlan
exit
!
interface Port-Channel3
description 00B_access_3
switchport mode trunk
switchport trunk allowed vlan add 2
switchport forbidden default-vlan
exit
!
interface vlan 1
shutdown
exit
!
interface vlan 2
ip address 192.168.50.1 255.255.255.0
ip dhcp relay enable
exit
!
interface vlan 111
ip address 192.168.11.10 255.255.255.0
exit
!
!
end
```

11.2.2 OOB access

```
no spanning-tree
!
vlan database
vlan 2
exit
!
hostname OOB_Access_1
!
interface gigabitethernet1/0/1
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/2
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/3
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/4
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/5
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/6
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/7
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/8
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/9
description OOB
switchport access vlan 2
exit
!
interface gigabitethernet1/0/10
description OOB
switchport access vlan 2
exit
```



```
!  
interface gigabitethernet1/0/11  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/12  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/13  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/14  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/15  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/16  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/17  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/18  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/19  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/20  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/21  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/22  
description 00B  
switchport access vlan 2  
exit
```

```
!  
interface gigabitethernet1/0/23  
description 00B  
switchport access vlan 2  
exit  
!  
interface gigabitethernet1/0/24  
description 00B  
switchport access vlan 2  
exit  
!  
interface tengigabitethernet1/0/1  
channel-group 1 mode auto  
exit  
!  
interface tengigabitethernet1/0/2  
channel-group 1 mode auto  
exit  
!  
interface Port-channel1  
description 00B_Core  
switchport mode trunk  
switchport trunk allowed vlan add 2  
switchport forbidden default-vlan  
exit  
!  
interface vlan 1  
shutdown  
exit  
!  
!  
end
```

12 Appendix 1

The Appendix contains full configurations of the devices used in this guide.

12.1 Configurations for IS-IS

12.1.1 Spine_1

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_1
!
interface HundredGigabitEthernet1/0/1
description Leaf_1
ip address 172.16.1.2 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/2
description Leaf_2
ip address 172.16.2.2 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/3
description Leaf_3
ip address 172.16.3.2 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface loopback1
ip address 10.0.1.1 255.255.255.255
exit
!
!
router bgp 65500
  bgp router-id 10.0.1.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
  exit
  !
```

```
    address-family l2vpn evpn
    exit
exit
!
neighbor 10.0.0.2
peer-group LEAF_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
!
neighbor 10.0.0.3
peer-group LEAF_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
router isis
address-family ipv4 unicast
redistribute connected
exit
net 49.0001.1111.1111.1111.00
exit
!
!
end
```

12.1.2 Spine_2

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_2
!
interface HundredGigabitEthernet1/0/1
description Leaf_1
ip address 172.16.1.6 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/2
description Leaf_2
ip address 172.16.2.6 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/3
description Leaf_3
ip address 172.16.3.6 255.255.255.252
ip router isis
isis network point-to-point
exit
!
interface loopback1
ip address 10.0.2.2 255.255.255.255
exit
!
!
router bgp 65500
  bgp router-id 10.0.2.2
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
```

```
!  
neighbor 10.0.0.2  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
!  
  address-family l2vpn evpn  
  exit  
exit  
!  
neighbor 10.0.0.3  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
!  
  address-family l2vpn evpn  
  exit  
exit  
exit  
!  
router isis  
  address-family ipv4 unicast  
  redistribute connected  
  exit  
  net 49.0001.2222.2222.2222.00  
  exit  
!  
!  
end
```

12.1.3 Leaf_1

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_1
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.1.1 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.1.5 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface loopback1
  ip address 10.0.0.1 255.255.255.255
exit
!
!
router bgp 65500
  bgp router-id 10.0.0.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
```



```
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
!
neighbor 10.0.2.2
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
router isis
address-family ipv4 unicast
redistribute connected
exit
net 49.0001.0001.0001.0001.00
exit
!
!
end
```

12.1.4 Leaf_2

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_2
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.2.1 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.2.5 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface loopback1
  ip address 10.0.0.2 255.255.255.255
exit
!
!
router bgp 65500
  bgp router-id 10.0.0.2
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
```

```
    exit
exit
!
neighbor 10.0.2.2
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
router isis
  address-family ipv4 unicast
  redistribute connected
  exit
  net 49.0001.0002.0002.0002.00
exit
!
!
end
```

12.1.5 Leaf_3

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.3.5 255.255.255.252
  ip router isis
  isis network point-to-point
exit
!
interface loopback1
  ip address 10.0.0.3 255.255.255.255
exit
!
!
router bgp 65500
  bgp router-id 10.0.0.3
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
```

```
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
!
neighbor 10.0.2.2
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
router isis
address-family ipv4 unicast
redistribute connected
exit
net 49.0001.0003.0003.0003.00
exit
!
!
end
```

12.2 Configurations for OSPF

12.2.1 Spine_1

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_1
!
line console
  exec-timeout 0
exit
!
interface HundredGigabitEthernet1/0/1
  description Leaf_1
  ip address 172.16.1.2 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Leaf_2
  ip address 172.16.2.2 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/3
  description Leaf_3
  ip address 172.16.3.2 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.1.1 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.1.2 area 0.0.0.0
  network 172.16.2.2 area 0.0.0.0
  network 172.16.3.2 area 0.0.0.0
  router-id 10.0.1.1
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.2
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.2
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.2
  ip ospf network point-to-point
exit
!
router bgp 65500
```

```
bgp router-id 10.0.1.1
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
!
peer-group LEAF_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd
  route-reflector-client
exit
!
neighbor 10.0.0.1
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.0.2
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.0.3
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
!
end
```

12.2.2 Spine_2

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_2
!
line console
  exec-timeout 0
exit
!
interface HundredGigabitEthernet1/0/1
  description Leaf_1
  ip address 172.16.1.6 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Leaf_2
  ip address 172.16.2.6 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/3
  description Leaf_3
  ip address 172.16.3.6 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.2.2 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.1.6 area 0.0.0.0
  network 172.16.2.6 area 0.0.0.0
  network 172.16.3.6 area 0.0.0.0
  router-id 10.0.2.2
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.6
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.2.2
  address-family ipv4 unicast
  exit
```



```
!  
address-family l2vpn evpn  
exit  
!  
peer-group LEAF_GROUP  
  remote-as 65500  
  update-source loopback 1  
  fall-over bfd  
  route-reflector-client  
exit  
!  
neighbor 10.0.0.1  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
!  
neighbor 10.0.0.2  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
!  
neighbor 10.0.0.3  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
exit  
!  
!  
end
```

12.2.3 Leaf_1

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_1
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.1.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.1.5 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.0.1 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.1.1 area 0.0.0.0
  network 172.16.1.5 area 0.0.0.0
  router-id 10.0.0.1
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.1.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.1
  address-family ipv4 unicast
```

```
exit
!
address-family l2vpn evpn
exit
!
peer-group SPINE_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd
exit
!
neighbor 10.0.1.1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.2.2
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
!
end
```

12.2.4 Leaf_2

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_2
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.2.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.2.5 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.0.2 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.2.1 area 0.0.0.0
  network 172.16.2.5 area 0.0.0.0
  router-id 10.0.0.2
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.2.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.2
  address-family ipv4 unicast
  exit
!
  address-family l2vpn evpn
  exit
!
```

```
peer-group SPINE_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd
exit
!
neighbor 10.0.1.1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.2.2
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
!
end
```

12.2.5 Leaf_3

```
no spanning-tree
!
vlan database
  vlan 1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
interface TenGigabitEthernet1/0/11
  description Host2
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.3.5 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.0.3 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.3.1 area 0.0.0.0
  network 172.16.3.5 area 0.0.0.0
  router-id 10.0.0.3
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.3.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.3
  address-family ipv4 unicast
```

```
exit
!
address-family l2vpn evpn
exit
!
peer-group SPINE_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd
exit
!
neighbor 10.0.1.1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.2.2
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
!
end
```

12.3 Configurations for Multicast VXLAN

12.3.1 Spine_1

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_1
!
interface HundredGigabitEthernet1/0/1
description Leaf_1
ip address 172.16.1.2 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/2
description Leaf_2
ip address 172.16.2.2 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/3
description Leaf_3
ip address 172.16.3.2 255.255.255.252
ip pim
exit
!
interface loopback1
ip address 10.0.1.1 255.255.255.255
exit
!
interface loopback2
ip address 10.100.100.100 255.255.255.255
description Anycast_RP_IP
exit
!
!
router ospf 1
network 172.16.1.2 area 0.0.0.0
network 172.16.2.2 area 0.0.0.0
network 172.16.3.2 area 0.0.0.0
router-id 10.0.1.1
timers spf delay 0
redistribute connected subnets
exit
!
interface ip 172.16.1.2
ip ospf network point-to-point
exit
!
interface ip 172.16.2.2
ip ospf network point-to-point
exit
!
interface ip 172.16.3.2
```



```

ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.1.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.2
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.3
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
router msdp
  connect-source 10.0.1.1
  originator-ip 10.100.100.100
  !
  peer 10.0.2.2
    mesh-group TESTGR
  exit
exit
!

```

end

12.3.2 Spine_2

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_2
!
interface HundredGigabitEthernet1/0/1
  description Leaf_1
  ip address 172.16.1.6 255.255.255.252
  ip pim
exit
!
interface HundredGigabitEthernet1/0/2
  description Leaf_2
  ip address 172.16.2.6 255.255.255.252
  ip pim
exit
!
interface HundredGigabitEthernet1/0/3
  description Leaf_3
  ip address 172.16.3.6 255.255.255.252
  ip pim
exit
!
interface loopback1
  ip address 10.0.2.2 255.255.255.255
exit
!
interface loopback2
  ip address 10.100.100.100 255.255.255.255
  description Anycast_RP_IP
exit
!
!
router ospf 1
  network 172.16.1.6 area 0.0.0.0
  network 172.16.2.6 area 0.0.0.0
  network 172.16.3.6 area 0.0.0.0
  router-id 10.0.2.2
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.6
  ip ospf network point-to-point
exit
!
```

```

router bgp 65500
  bgp router-id 10.0.2.2
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.2
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.3
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
router msdp
  connect-source 10.0.2.2
  originator-ip 10.100.100.100
  !
  peer 10.0.1.1
    mesh-group TESTGR
  exit
exit
!
end

```

12.3.3 Leaf_1

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_1
!
interface TenGigabitEthernet1/0/9
  description Host1_mcast
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  switchport forbidden default-vlan
exit
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.1.1 255.255.255.252
```

```

    ip pim
exit
!
interface HundredGigabitEthernet1/0/2
    description Spine_2
    ip address 172.16.1.5 255.255.255.252
    ip pim
exit
!
interface loopback1
    ip address 10.0.0.1 255.255.255.255
    ip pim
    ip igmp static-group 233.0.0.2
    ip igmp static-group 233.0.0.3
    ip igmp static-group 233.0.0.4
    ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
    network 172.16.1.1 area 0.0.0.0
    network 172.16.1.5 area 0.0.0.0
    router-id 10.0.0.1
    timers spf delay 0
    redistribute connected subnets
exit
!
interface ip 172.16.1.1
    ip ospf network point-to-point
exit
!
interface ip 172.16.1.5
    ip ospf network point-to-point
exit
!
router bgp 65500
    bgp router-id 10.0.0.1
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
    !
    peer-group SPINE_GROUP
        remote-as 65500
        update-source loopback 1
        fall-over bfd
    exit
    !
    neighbor 10.0.1.1
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
    neighbor 10.0.2.2
        peer-group SPINE_GROUP

```

```
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
ip multicast multipath group-paths-num
!
end
```

12.3.4 Leaf_2

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_2
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.2.1 255.255.255.252
  ip pim
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.2.5 255.255.255.252
  ip pim
exit
!
interface loopback1
  ip address 10.0.0.2 255.255.255.255
  ip pim
```



```

ip igmp static-group 233.0.0.2
ip igmp static-group 233.0.0.3
ip igmp static-group 233.0.0.4
ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
network 172.16.2.1 area 0.0.0.0
network 172.16.2.5 area 0.0.0.0
router-id 10.0.0.2
timers spf delay 0
redistribute connected subnets
exit
!
interface ip 172.16.2.1
ip ospf network point-to-point
exit
!
interface ip 172.16.2.5
ip ospf network point-to-point
exit
!
router bgp 65500
bgp router-id 10.0.0.2
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
!
peer-group SPINE_GROUP
remote-as 65500
update-source loopback 1
fall-over bfd
exit
!
neighbor 10.0.1.1
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
!
neighbor 10.0.2.2
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100

```

```
!  
ip multicast multipath group-paths-num  
!  
end
```

12.3.5 Leaf_3

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
interface TenGigabitEthernet1/0/9
  description Host2_mcast
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  switchport forbidden default-vlan
exit
!
interface TenGigabitEthernet1/0/11
  description Host2
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
```

```

    ip pim
exit
!
interface HundredGigabitEthernet1/0/2
    description Spine_2
    ip address 172.16.3.5 255.255.255.252
    ip pim
exit
!
interface loopback1
    ip address 10.0.0.3 255.255.255.255
    ip pim
    ip igmp static-group 233.0.0.2
    ip igmp static-group 233.0.0.3
    ip igmp static-group 233.0.0.4
    ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
    network 172.16.3.1 area 0.0.0.0
    network 172.16.3.5 area 0.0.0.0
    router-id 10.0.0.3
    timers spf delay 0
    redistribute connected subnets
exit
!
interface ip 172.16.3.1
    ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
    ip ospf network point-to-point
exit
!
router bgp 65500
    bgp router-id 10.0.0.3
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
    !
    peer-group SPINE_GROUP
        remote-as 65500
        update-source loopback 1
        fall-over bfd
    exit
    !
    neighbor 10.0.1.1
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
    neighbor 10.0.2.2
        peer-group SPINE_GROUP

```

```
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
ip multicast multipath group-paths-num
!
end
```

12.4 Configurations for Symmetric IRB

12.4.1 Spine_1

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_1
!
line console
  exec-timeout 0
exit
!
interface HundredGigabitEthernet1/0/1
  description Leaf_1
  ip address 172.16.1.2 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Leaf_2
  ip address 172.16.2.2 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/3
  description Leaf_3
  ip address 172.16.3.2 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.1.1 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.1.2 area 0.0.0.0
  network 172.16.2.2 area 0.0.0.0
  network 172.16.3.2 area 0.0.0.0
  router-id 10.0.1.1
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.2
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.2
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.2
  ip ospf network point-to-point
exit
!
router bgp 65500
```

```
bgp router-id 10.0.1.1
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
!
peer-group LEAF_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd
  route-reflector-client
exit
!
neighbor 10.0.0.1
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.0.2
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
neighbor 10.0.0.3
  peer-group LEAF_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
exit
!
!
end
```

12.4.2 Spine_2

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_2
!
line console
  exec-timeout 0
exit
!
interface HundredGigabitEthernet1/0/1
  description Leaf_1
  ip address 172.16.1.6 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Leaf_2
  ip address 172.16.2.6 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/3
  description Leaf_3
  ip address 172.16.3.6 255.255.255.252
exit
!
interface loopback1
  ip address 10.0.2.2 255.255.255.255
exit
!
!
router ospf 1
  network 172.16.1.6 area 0.0.0.0
  network 172.16.2.6 area 0.0.0.0
  network 172.16.3.6 area 0.0.0.0
  router-id 10.0.2.2
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.6
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.6
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.2.2
  address-family ipv4 unicast
  exit
```



```
!  
address-family l2vpn evpn  
exit  
!  
peer-group LEAF_GROUP  
  remote-as 65500  
  update-source loopback 1  
  fall-over bfd  
  route-reflector-client  
exit  
!  
neighbor 10.0.0.1  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
!  
neighbor 10.0.0.2  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
!  
neighbor 10.0.0.3  
  peer-group LEAF_GROUP  
  address-family ipv4 unicast  
  exit  
  !  
  address-family l2vpn evpn  
  exit  
exit  
exit  
!  
!  
end
```

12.4.3 Leaf_1

```
ip vrf VRF1
  vni 100100
  route-target both 65500:100100
exit
!
!
no spanning-tree
!
vlan database
  vlan 100,1000
exit
!
port jumbo-frame
!
vxlan L3_vxlan
  vni 100100 ip-routing
  vlan 100
exit
!
vxlan test_vxlan1
  vni 101000
  vlan 1000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_1
!
line console
  exec-timeout 0
exit
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.1.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.1.5 255.255.255.252
exit
!
interface vlan 100
  ip vrf VRF1
exit
!
interface vlan 1000
  ip vrf VRF1
  ip address 192.168.1.1 255.255.255.0
exit
!
```

```

interface loopback1
  ip address 10.0.0.1 255.255.255.255
exit
!
!
!
!
router ospf 1
  network 172.16.1.1 area 0.0.0.0
  network 172.16.1.5 area 0.0.0.0
  router-id 10.0.0.1
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.1.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.2.2
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
vrf VRF1
  address-family ipv4 unicast
    redistribute connected
  exit
exit

```

```
exit  
!  
!  
end
```

12.4.4 Leaf_2

```
ip vrf VRF1
  vni 100100
  route-target both 65500:100100
exit
!
!
no spanning-tree
!
vlan database
  vlan 100,2000
exit
!
port jumbo-frame
!
vxlan L3_vxlan
  vni 100100 ip-routing
  vlan 100
exit
!
vxlan test_vxlan2
  vni 102000
  vlan 2000
exit
!
!
ip maximum-paths 32
!
anycast-gateway mac-address 00:00:00:11:11:11
!
hostname Leaf_2
!
line console
  exec-timeout 0
exit
!
interface TenGigabitEthernet1/0/11
  description Host2
  switchport access vlan 2000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.2.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.2.5 255.255.255.252
exit
!
interface vlan 100
  ip vrf VRF1
exit
!
interface vlan 2000
  ip vrf VRF1
  ip address 192.168.2.1 255.255.255.0
```

```

    anycast-gateway
exit
!
interface loopback1
    ip address 10.0.0.2 255.255.255.255
exit
!
!
!
!
router ospf 1
    network 172.16.2.1 area 0.0.0.0
    network 172.16.2.5 area 0.0.0.0
    router-id 10.0.0.2
    timers spf delay 0
    redistribute connected subnets
exit
!
interface ip 172.16.2.1
    ip ospf network point-to-point
exit
!
interface ip 172.16.2.5
    ip ospf network point-to-point
exit
!
router bgp 65500
    bgp router-id 10.0.0.2
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
    !
    peer-group SPINE_GROUP
        remote-as 65500
        update-source loopback 1
        fall-over bfd
    exit
    !
    neighbor 10.0.1.1
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
    neighbor 10.0.2.2
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
vrf VRF1
    address-family ipv4 unicast

```

```
    redistribute connected
  exit
exit
exit
!
!
end
```

12.4.5 Leaf_3

```
ip vrf VRF1
  vni 100100
  route-target both 65500:100100
exit
!
!
no spanning-tree
!
vlan database
  vlan 100,2000
exit
!
port jumbo-frame
!
vxlan L3_vxlan
  vni 100100 ip-routing
  vlan 100
exit
!
vxlan test_vxlan2
  vni 102000
  vlan 2000
exit
!
!
ip maximum-paths 32
!
anycast-gateway mac-address 00:00:00:11:11:11
!
hostname Leaf_3
!
interface TenGigabitEthernet1/0/11
  description Host3
  switchport access vlan 2000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.3.5 255.255.255.252
exit
!
interface vlan 100
  ip vrf VRF1
exit
!
interface vlan 2000
  ip vrf VRF1
  ip address 192.168.2.1 255.255.255.0
  anycast-gateway
exit
!
interface loopback1
```



```

    ip address 10.0.0.3 255.255.255.255
exit
!
!
!
!
router ospf 1
    network 172.16.3.1 area 0.0.0.0
    network 172.16.3.5 area 0.0.0.0
    router-id 10.0.0.3
    timers spf delay 0
    redistribute connected subnets
exit
!
interface ip 172.16.3.1
    ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
    ip ospf network point-to-point
exit
!
router bgp 65500
    bgp router-id 10.0.0.3
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
    !
    peer-group SPINE_GROUP
        remote-as 65500
        update-source loopback 1
        fall-over bfd
    exit
    !
    neighbor 10.0.1.1
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
    neighbor 10.0.2.2
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
vrf VRF1
    address-family ipv4 unicast
        redistribute connected
    exit
exit
exit

```

!
!
end

12.4.6 Leaf_3 + OSPF

```
ip vrf VRF1
  vni 100100
  route-target both 65500:100100
exit
!
!
no spanning-tree
!
vlan database
  vlan 100,2000
exit
!
port jumbo-frame
!
vxlan L3_vxlan
  vni 100100 ip-routing
  vlan 100
exit
!
vxlan test_vxlan2
  vni 102000
  vlan 2000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
line console
  exec-timeout 0
exit
!
interface TenGigabitEthernet1/0/11
  description Host3
  switchport access vlan 2000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.3.5 255.255.255.252
exit
!
interface vlan 100
  ip vrf VRF1
exit
!
interface vlan 2000
  ip vrf VRF1
  ip address 192.168.2.1 255.255.255.0
exit
!
```

```

interface loopback1
  ip address 10.0.0.3 255.255.255.255
exit
!
!
!
!
router ospf 1
  network 172.16.3.1 area 0.0.0.0
  network 172.16.3.5 area 0.0.0.0
  router-id 10.0.0.3
  timers spf delay 0
  redistribute connected subnets
exit
!
router ospf 2 vrf VRF1
  network 192.168.2.1 area 10.10.10.10
  router-id 192.168.2.1
  redistribute bgp subnets
exit
!
interface ip 172.16.3.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.3
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.2.2
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit

```

```
!  
vrf VRF1  
  address-family ipv4 unicast  
    redistribute connected  
    redistribute ospf 2 match internal  
    redistribute ospf 2 match external-1  
    redistribute ospf 2 match external-2  
  exit  
exit  
exit  
!  
!  
end
```

12.4.7 Leaf_3 + eBGP

```
ip vrf VRF1
  vni 100100
  route-target both 65500:100100
exit
!
!
no spanning-tree
!
vlan database
  vlan 100,2000
exit
!
port jumbo-frame
!
vxlan L3_vxlan
  vni 100100 ip-routing
  vlan 100
exit
!
vxlan test_vxlan2
  vni 102000
  vlan 2000
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
line console
  exec-timeout 0
exit
!
interface TenGigabitEthernet1/0/11
  description Host3
  switchport access vlan 2000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.3.5 255.255.255.252
exit
!
interface vlan 100
  ip vrf VRF1
exit
!
interface vlan 2000
  ip vrf VRF1
  ip address 192.168.2.1 255.255.255.0
exit
!
```

```

interface loopback1
  ip address 10.0.0.3 255.255.255.255
exit
!
!
!
!
router ospf 1
  network 172.16.3.1 area 0.0.0.0
  network 172.16.3.5 area 0.0.0.0
  router-id 10.0.0.3
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.3.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.3
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.2.2
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
vrf VRF1
  address-family ipv4 unicast
    redistribute connected
    redistribute bgp
  exit

```

```
!  
neighbor 192.168.2.2  
  remote-as 65600  
  address-family ipv4 unicast  
  exit  
exit  
exit  
exit  
!  
!  
end
```


12.5 Configurations for EVPN multihoming

12.5.1 Spine_1

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_1
!
interface HundredGigabitEthernet1/0/1
description Leaf_1
ip address 172.16.1.2 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/2
description Leaf_2
ip address 172.16.2.2 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/3
description Leaf_3
ip address 172.16.3.2 255.255.255.252
ip pim
exit
!
interface loopback1
ip address 10.0.1.1 255.255.255.255
exit
!
interface loopback2
ip address 10.100.100.100 255.255.255.255
description Anycast_RP_IP
exit
!
!
router ospf 1
network 172.16.1.2 area 0.0.0.0
network 172.16.2.2 area 0.0.0.0
network 172.16.3.2 area 0.0.0.0
router-id 10.0.1.1
timers spf delay 0
redistribute connected subnets
exit
!
interface ip 172.16.1.2
ip ospf network point-to-point
exit
!
interface ip 172.16.2.2
ip ospf network point-to-point
exit
!
interface ip 172.16.3.2
```

```

ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.1.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.2
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.3
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
router msdp
  connect-source 10.0.1.1
  originator-ip 10.100.100.100
  !
  peer 10.0.2.2
    mesh-group TESTGR
  exit
exit
!

```

end

12.5.2 Spine_2

```
no spanning-tree
!
port jumbo-frame
!
ip maximum-paths 32
!
hostname Spine_2
!
interface HundredGigabitEthernet1/0/1
description Leaf_1
ip address 172.16.1.6 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/2
description Leaf_2
ip address 172.16.2.6 255.255.255.252
ip pim
exit
!
interface HundredGigabitEthernet1/0/3
description Leaf_3
ip address 172.16.3.6 255.255.255.252
ip pim
exit
!
interface loopback1
ip address 10.0.2.2 255.255.255.255
exit
!
interface loopback2
ip address 10.100.100.100 255.255.255.255
description Anycast_RP_IP
exit
!
!
router ospf 1
network 172.16.1.6 area 0.0.0.0
network 172.16.2.6 area 0.0.0.0
network 172.16.3.6 area 0.0.0.0
router-id 10.0.2.2
timers spf delay 0
redistribute connected subnets
exit
!
interface ip 172.16.1.6
ip ospf network point-to-point
exit
!
interface ip 172.16.2.6
ip ospf network point-to-point
exit
!
interface ip 172.16.3.6
ip ospf network point-to-point
exit
!
```

```

router bgp 65500
  bgp router-id 10.0.2.2
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group LEAF_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
    route-reflector-client
  exit
  !
  neighbor 10.0.0.1
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.2
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
  !
  neighbor 10.0.0.3
    peer-group LEAF_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
  exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
router msdp
  connect-source 10.0.2.2
  originator-ip 10.100.100.100
  !
  peer 10.0.1.1
    mesh-group TESTGR
  exit
exit
!
end

```

12.5.3 Leaf_1

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_1
!
interface TenGigabitEthernet1/0/9
  description Host1_mcast
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  switchport forbidden default-vlan
exit
!
interface TenGigabitEthernet1/0/11
  description Host1
  switchport access vlan 1000
exit
!
interface TenGigabitEthernet1/0/12
  channel-group 1 mode auto
exit
```

```

!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.1.1 255.255.255.252
  ip pim
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.1.5 255.255.255.252
  ip pim
exit
!
interface Port-Channel1
  description Host1_LAG
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  ethernet-segment 1000
  system-mac 11:22:33:44:55:66
  exit
exit
!
interface loopback1
  ip address 10.0.0.1 255.255.255.255
  ip pim
  ip igmp static-group 233.0.0.2
  ip igmp static-group 233.0.0.3
  ip igmp static-group 233.0.0.4
  ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
  network 172.16.1.1 area 0.0.0.0
  network 172.16.1.5 area 0.0.0.0
  router-id 10.0.0.1
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.1.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.1.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.1
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
  remote-as 65500
  update-source loopback 1
  fall-over bfd

```

```
exit
!
neighbor 10.0.1.1
  advertisement-interval 1 withdraw 1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
!
neighbor 10.0.2.2
  advertisement-interval 1 withdraw 1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
ip multicast multipath group-paths-num
!
end
```


12.5.4 Leaf_2

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_2
!
interface TenGigabitEthernet1/0/12
  channel-group 1 mode auto
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.2.1 255.255.255.252
  ip pim
exit
!
interface HundredGigabitEthernet1/0/2
  description Spine_2
  ip address 172.16.2.5 255.255.255.252
  ip pim
exit
```

```

!
interface Port-Channel1
  description Host1_LAG
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  ethernet-segment 1000
    system-mac 11:22:33:44:55:66
  exit
exit
!
interface loopback1
  ip address 10.0.0.2 255.255.255.255
  ip pim
  ip igmp static-group 233.0.0.2
  ip igmp static-group 233.0.0.3
  ip igmp static-group 233.0.0.4
  ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
  network 172.16.2.1 area 0.0.0.0
  network 172.16.2.5 area 0.0.0.0
  router-id 10.0.0.2
  timers spf delay 0
  redistribute connected subnets
exit
!
interface ip 172.16.2.1
  ip ospf network point-to-point
exit
!
interface ip 172.16.2.5
  ip ospf network point-to-point
exit
!
router bgp 65500
  bgp router-id 10.0.0.2
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
  !
  peer-group SPINE_GROUP
    remote-as 65500
    update-source loopback 1
    fall-over bfd
  exit
  !
  neighbor 10.0.1.1
    advertisement-interval 1 withdraw 1
    peer-group SPINE_GROUP
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
exit
!

```

```
neighbor 10.0.2.2
  advertisement-interval 1 withdraw 1
  peer-group SPINE_GROUP
  address-family ipv4 unicast
  exit
  !
  address-family l2vpn evpn
  exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
ip multicast multipath group-paths-num
!
end
```

12.5.5 Leaf_3

```
no spanning-tree
!
vlan database
  vlan 2-5,1000
exit
!
port jumbo-frame
!
vxlan test_vxlan
  vni 101000
  vlan 1000
exit
!
vxlan mcast2
  vni 102
  vlan 2
  mcast-group 233.0.0.2
exit
!
vxlan mcast3
  vni 103
  vlan 3
  mcast-group 233.0.0.3
exit
!
vxlan mcast4
  vni 104
  vlan 4
  mcast-group 233.0.0.4
exit
!
vxlan mcast5
  vni 105
  vlan 5
  mcast-group 233.0.0.5
exit
!
!
ip maximum-paths 32
!
hostname Leaf_3
!
interface TenGigabitEthernet1/0/9
  description Host2_mcast
  switchport mode trunk
  switchport trunk allowed vlan add 2-5
  switchport forbidden default-vlan
exit
!
interface TenGigabitEthernet1/0/11
  description Host2
  switchport access vlan 1000
exit
!
interface HundredGigabitEthernet1/0/1
  description Spine_1
  ip address 172.16.3.1 255.255.255.252
```

```

    ip pim
exit
!
interface HundredGigabitEthernet1/0/2
    description Spine_2
    ip address 172.16.3.5 255.255.255.252
    ip pim
exit
!
interface loopback1
    ip address 10.0.0.3 255.255.255.255
    ip pim
    ip igmp static-group 233.0.0.2
    ip igmp static-group 233.0.0.3
    ip igmp static-group 233.0.0.4
    ip igmp static-group 233.0.0.5
exit
!
!
router ospf 1
    network 172.16.3.1 area 0.0.0.0
    network 172.16.3.5 area 0.0.0.0
    router-id 10.0.0.3
    timers spf delay 0
    redistribute connected subnets
exit
!
interface ip 172.16.3.1
    ip ospf network point-to-point
exit
!
interface ip 172.16.3.5
    ip ospf network point-to-point
exit
!
router bgp 65500
    bgp router-id 10.0.0.3
    address-family ipv4 unicast
    exit
    !
    address-family l2vpn evpn
    exit
    !
    peer-group SPINE_GROUP
        remote-as 65500
        update-source loopback 1
        fall-over bfd
    exit
    !
    neighbor 10.0.1.1
        advertisement-interval 1 withdraw 1
        peer-group SPINE_GROUP
        address-family ipv4 unicast
        exit
        !
        address-family l2vpn evpn
        exit
    exit
    !
    neighbor 10.0.2.2

```

```
advertisement-interval 1 withdraw 1
peer-group SPINE_GROUP
address-family ipv4 unicast
exit
!
address-family l2vpn evpn
exit
exit
exit
!
!
ip multicast-routing pim
!
ip pim rp-address 10.100.100.100
!
ip multicast multipath group-paths-num
!
end
```

TECHNICAL SUPPORT

For technical assistance in issues related to handling Eltex Ltd. equipment, please, address to Service Center of the company:

<http://www.eltex-co.com/support>

You are welcome to visit Eltex official website to get the relevant technical documentation and software, to use our knowledge base or consult a Service Center Specialist in our technical forum.

<http://www.eltex-co.com/>

<http://www.eltex-co.com/support/downloads/>