aruba
a Hewlett Packard Enterprise company

VALIDATED REFERENCE DESIGN GUIDE

AOS-CX 10.3 MULTICAST GUIDE

CONTENTS

INTRODUCTION

This guide provides information on IP multicast networking capabilities, use cases, and configuration best practices with Aruba AOS-CX switching platforms running code version 10.3. The tested scenario used in this document were focused on Data Center Network Spine-Leaf design, however, many of the technologies described here can be leveraged and deployed in Campus networks also. The intended audience for this document is IT administrators and solution architects planning on deploying IP Multicast features.

WHAT IS MULTICAST NETWORKING?

In computer networking, IP multicast is method used to enable data transmission to be addressed to a group of destination computers at the same time.

By enabling high-efficiency point-to-multipoint data transmission over a network, multicast is able to greatly save network bandwidth and reduces the network load. By using multicast technology, a network operator can easily provide bandwidth-critical and time-critical information services. These services include live webcasting, Web TV, distance learning, telemedicine, Web radio, and real-time video conferencing.

When some hosts on the network need multicast information, the multicast source sends only one copy of the information. On the network, Multicast distribution trees are built through multicast routing protocols, and the packets are replicated only on the nodes where the trees branch.

In the image below, the multicast source sends a single copy to the first hop switch. The switches and routers in the group will forward the flow to devices that only have hosts (receivers) that are interested in receiving the stream.

Figure 1. Multicast Networking

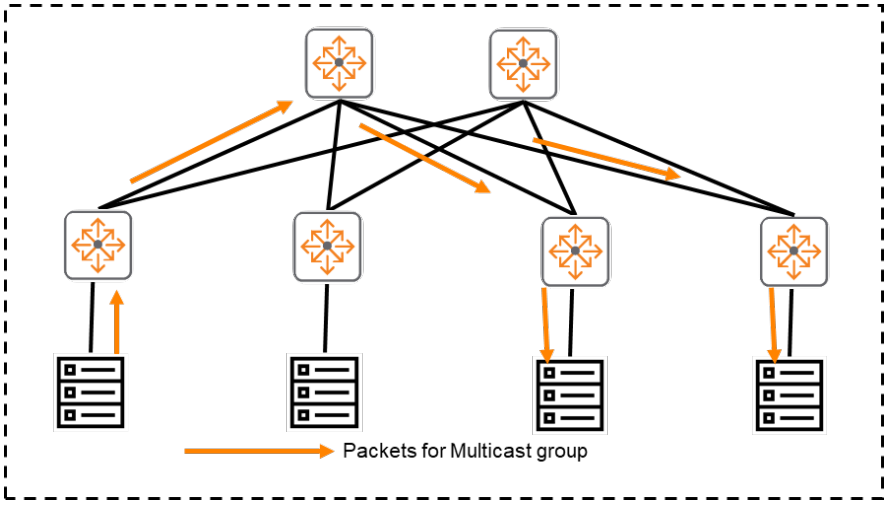

To summarize, multicast has the following advantages:

- Advantages over unicast—Multicast data is replicated and distributed until it flows to the farthest-possible node from the source. The increase of receiver hosts will not remarkably increase the load of the source or the usage of network resources.
- Advantages over broadcast—Multicast data is sent only to the receivers that need it. This saves network

bandwidth and enhances network security. In addition, multicast data is not confined to the same subnet.

MULTICAST PROTOCOLS

Protocol Independent Multicast (PIM)

PIM is a family of routing protocols that form multicast trees to forward traffic from multicast sources to subnets that have used a protocol such as IGMP to request the traffic. PIM relies on the unicast routing tables created by any of several unicast routing protocols to identify the path back to a multicast source (reverse path forwarding, or RPF). With this information, PIM sets up the distribution tree for the multicast traffic.

In a PIM-SM network, a Rendezvous Point (RP) acts as a central source of information for sources and receivers of multicast data. Sources send traffic to the RP which then forwarded the MC traffic to receivers down a distribution tree.

AOS-CX switches support PIM-Sparse Mode (PIM-SM). PIM-SM assumes that most hosts do not want to receive multicast traffic, and it uses a non-flooding multicast model to direct traffic for a particular multicast group from the source to the interface where there are multicast receivers that have joined the group. As a result, this model sends traffic only to the routers that have specifically requested it.

In this model if the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source

IGMP

Internet Group Management Protocol (IGMP) establishes and maintains the multicast group memberships between a Layer 3 multicast device and the subnet with directly connected hosts. The AOS-CX switches support IGMPv1, IGMPv2, and IGMPv3.

Additionally, IGMP snooping runs on the AOS-CX switches as a multicast constraining mechanism to help improve multicast forwarding efficiency. It creates Layer 2 multicast forwarding entries from the IGMP packets that are exchanged between the hosts and the router. When IGMP snooping is not enabled, the Layer 2 switch floods multicast packets to all hosts in a VLAN or VSI. When IGMP snooping is enabled, the Layer 2 switch forwards multicast packets of known multicast groups to only the receivers within that subnet.

Multicast Source Discovery Protocol (MSDP)

MSDP is used to connect multicast routing domains. It typically runs on the same router as PIM-SM rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

MULTICAST CONFIGURATION EXAMPLE

The topology as shown in the Figure below and sample configurations are described in this section.

Figure 2. Multicast Network Example Scenario

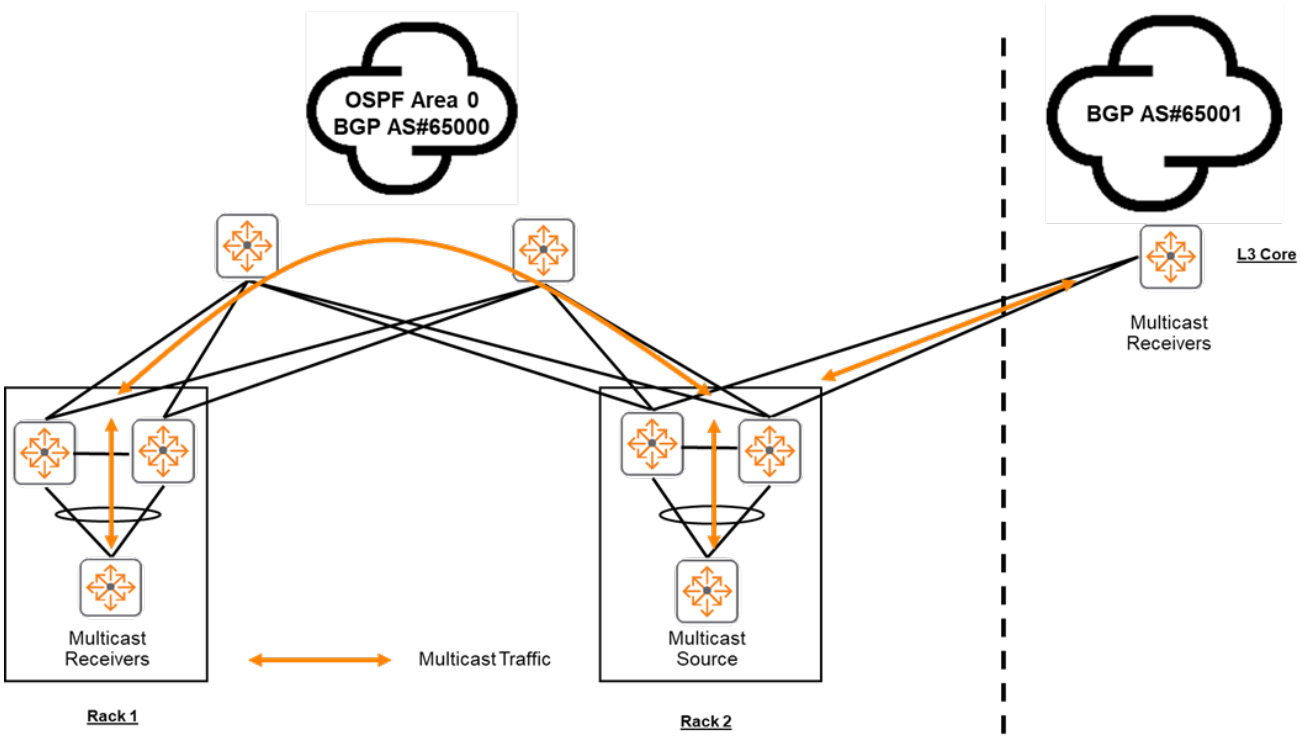

The above example solution is a Spine and Leaf solution with 2 racks - each rack contains 1 host and 2 Leaf switches. The Spine and Leaf solution shown uses OSPF as the Interior Gateway Protocol (IGP), and BGP to share routes to the other autonomous system (AS). This example solution also shows a device in another AS which is also interested in the multicast source streams.

Within the Spine and Leaf Pod the 2 Leaf switches in each rack are deployed in a highly available (HA) solution using Aruba VSX and Link Aggregation on the interfaces attached to the hosts. Each physical Leaf switch has an uplink connection to each Spine switch and uses OSPF as the Interior Gateway Protocol to share routes. Equal Cost Multipathing (ECMP) is used to load-balance unicast traffic across the Spine switches.

IGMP Snooping is configured on the VLANs and interfaces facing the attached hosts to ensure that those segments are listed as having hosts that are interested in the multicast stream. PIM-SM is used to share those multicast streams across the L3 OSPF fabric.

The Spine and Leaf Pod also connects to a L3 Core in a different AS. BGP is used for Route peering between these 2 AS numbers and Multicast Source Discovery Protocol (MSDP) is used to share Multicast PIM tree information for each PIM domain.

In the tests shown in this document, Rack 1 and the L3Core each have a host which is interested in the multicast source. Rack 2 connects to a host which is acting as the source for 5 multicast address groups.

Configuration recommendations and guidelines

- PIM-Sparse Mode (PIM-SM) must be configured on L3 physical interfaces, VLAN interfaces, and/or loopback interfaces that require Multicast traffic forwarding.
- The default PIM setting enables the spt-threshold setting. This ensures that the PIM domain will shift to the shortest path tree after learning the multicast traffic via the RPT.
- In 10.3:
    - BSR/Candidate-RP can only be used in a multicast domain without MSDP
    - As MSDP is used between different AS (#65000 and #65001), the Leafs in #65000 should utilize static RP
    - Do not create more than one MSDP session between two devices on the same VRF
- Ensure IGMP and IGMP Snooping are enabled on VLANs and VLAN interfaces facing host and clients that may be interested in receiving a multicast stream.
- When using VSX and IGMP Snooping, the switch with the lower SVI IP Address will become the IGMP querier to forward multicast traffic. The other VSX switch will become the IGMP proxy, waiting to forward multicast traffic in case the IGMP querier goes down.
- Within a VSX pair, one of the VSX switches will be the active forwarder. In case of switch failure, the PIM “active-active” configuration allows for faster failover to the other VSX switch.

CONFIGURATIONS DETAILS

Below are the configurations with details about the tested solutions.

The tested scenario used in this document is focused on Data Center Network Spine-Leaf design, however, many of the technologies described here can be leveraged and deployed in Campus networks also.

The next figure shows detailed solution.

Figure 3. Multicast Network Details

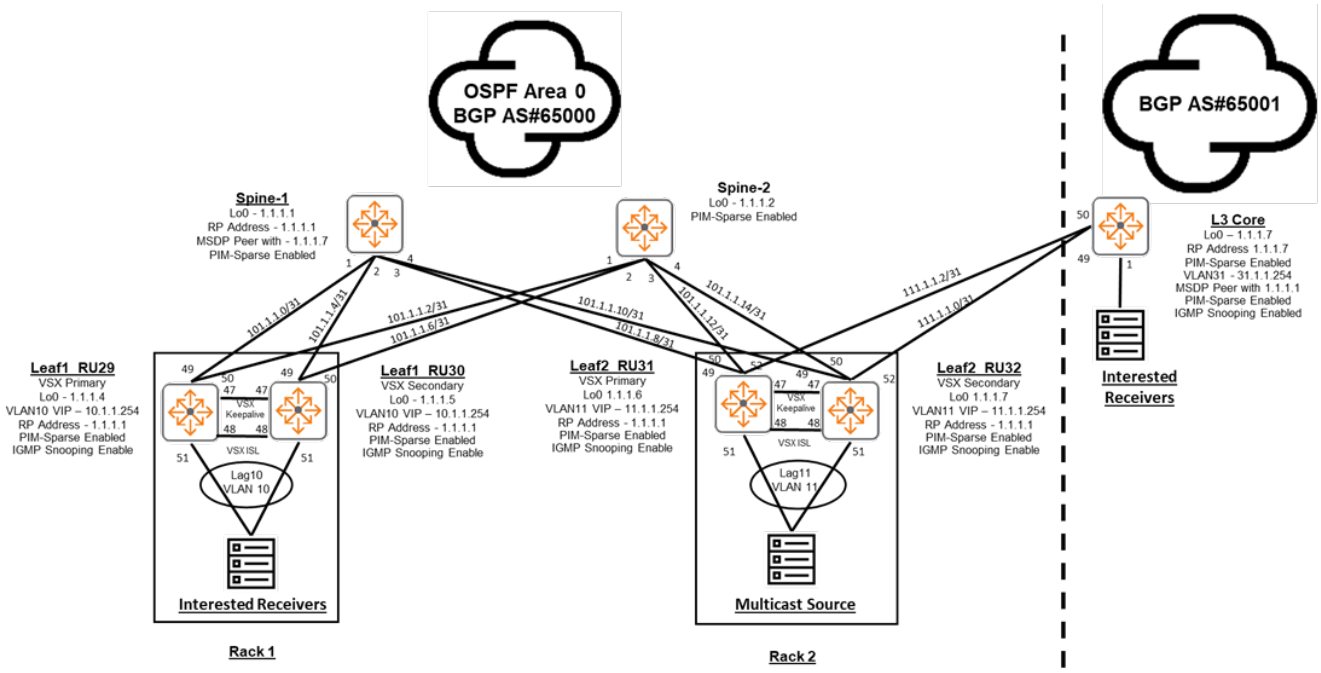

As shown in the image:

Rack 1 and Rack 2 (AS#65000)

- Rack 1 Leafs are connected to a single host.
  - Link Aggregation 10 is used to multi-chassis bundle interfaces 1/1/51 from both switches
  - LAG10 and the Host use VLAN 10, and the VSX switches are the default gateway
  - The Host in Rack 1 is interested in receiving multicast traffic
- Rack 2 Leafs are connected to a single host
  - Link Aggregation 11 is used to multi-chassis bundle interfaces 1/1/51 from both switches
  - LAG11 and the Host use VLAN 11, and the VSX switches are the default gateway
  - The Host in Rack 2 is the Multicast source
- The Leafs in each rack form a VSX pair using 1/1/48 as the ISL, and 1/1/47 as the Keepalive
- Each Leaf has a connection to each spine using ports 1/1/49 and 1/1/50
- /31 interfaces are used to provide direct IP connectivity between Leafs and Spines
- OSPF is used to share the rack to rack routes
- The Leaf VLANs and Loopbacks are distributed into the OSPF domain
- IGMP and IGMP Snooping is configured to ensure hosts in a segment can receive a multicast stream

- PIM-SM is used to ensure source multicast traffic can get routed to the interested receiving segments
- Rack 2 switches use BGP to share routes between the different AS numbers (65000 & 65001)
- Each rack points to Spine 1 (1.1.1.1) as the PIM RP

Spines (AS#65000)

- Each switch uses /31 interfaces on connections to each Leaf using ports 1/1/1 to 1/1/4
- Uses OSPF to share unicast routes between the racks
- Uses iBGP to peer with directly connected Leafs
- Configured as BGP Route Reflectors to allow iBGP peering between Leafs
- Spine 1 is configured as the RP
- PIM-SM is used to ensure source multicast traffic can get routed to the interested receiving segments
- MSDP is used to share Multicast RP information to share Multicast traffic between AS numbers 65000 & 65001

L3 Core (AS#65001)

- L3 Core device in the separate AS is connected to a single host
    - The Host uses VLAN 31
    - The Host is interested in receiving multicast traffic
- Uses /31 interfaces to connect to the Leafs in Rack 2 using ports 1/1/1 and 1/1/2
- Uses BGP to share routes between the different AS numbers 65000 & 65001
- The Leaf VLANs and Loopbacks are distributed into BGP domain
- IGMP and IGMP Snooping is configured to ensure hosts in a segment can receive a multicast stream
- PIM-SM is used to ensure source multicast traffic can get routed to the interested receiving segments
- MSDP is used to share Multicast RP information to share Multicast traffic between AS numbers 65000 & 65001
- A multicast generator is running and sending 5 multicast group addresses as available for interested listeners (239.0.5.1, 239.0.5.2, 239.0.5.3, 239.0.5.4, 239.0.5.5)

1) INITIAL SETUP (OOB, INITIAL CONFIGS)

In the DC, we recommend switches connect their management ports to an Out Of Band (OOB) management network as shown in Figure 2, this allows the switches to be manageable if there is an issue with In Band network connectivity.

Figure 2. OOB connectivity

Figure 4. OOB Network

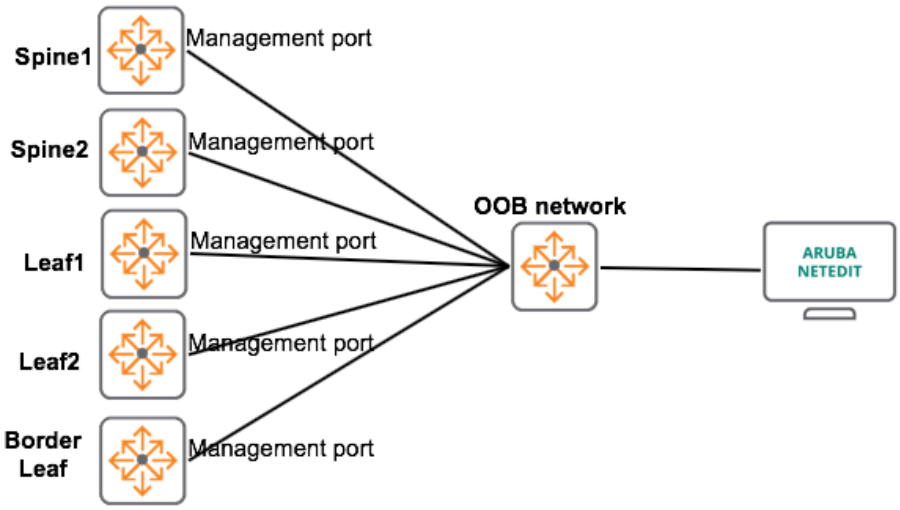

In order for NetEdit to manage each switch, initial configs should be added via one of these options:

- Aruba CX mobile app
- Console cable
- Zero Touch Provisioning (ZTP)

Sample Initial Configuration

```
hostname LEAF1
user admin group administrators password ciphertext AQBapUz+
!
!
ssh server vrf mgmt
!
!
!
!
interface mgmt
    no shutdown
    ip static 10.6.8.19/24
    default-gateway 10.6.8.1
!
! interface group 1 contains ports 1/1/1-1/1/12
system interface-group 1 speed 10g
! interface group 4 contains ports 1/1/37-1/1/48
system interface-group 4 speed 10g
!
https-server rest access-mode read-write
https-server vrf mgmt
```

Once the switches are configured and physically connected, ensure NetEdit has IP connectivity to switch management IPs and add all devices into NetEdit [Devices -> Action -> Add Device(s) or Add Multiple Devices]

Figure 5. Adding devices

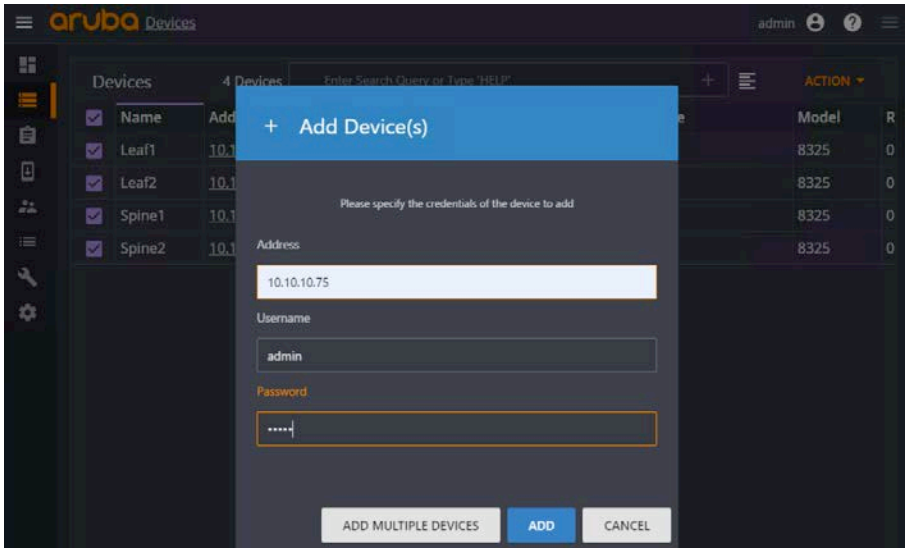

You should now see all managed devices.

Figure 6. Devices added

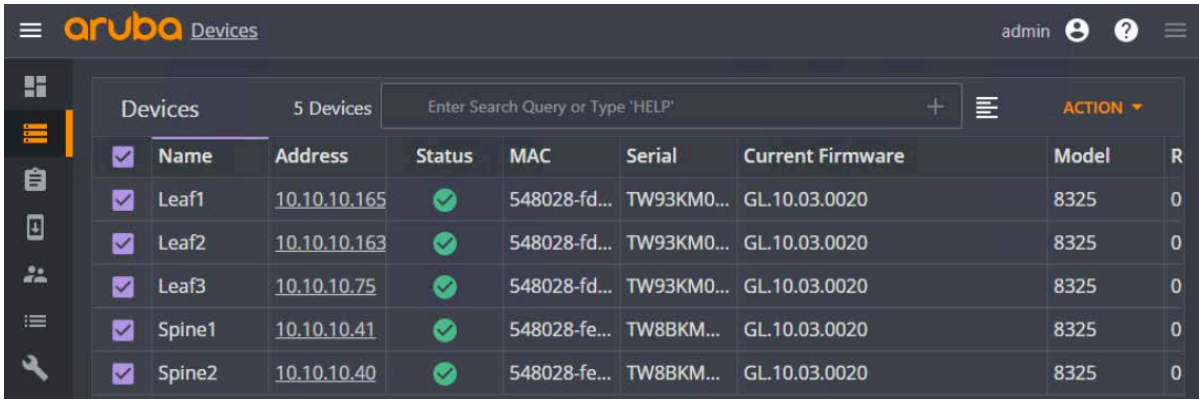

2) MODIFY CHANGE VALIDATION SETTINGS

You can add or modify the change validation commands used by NetEdit to enable or disable additional validation commands that could assist with specific deployments.

To help with available change validation captures,

This is done in NetEdit [Settings -> Validation -> Change Validation -> Command Scripts]

Figure 7. Adding additional commands to Change Validation

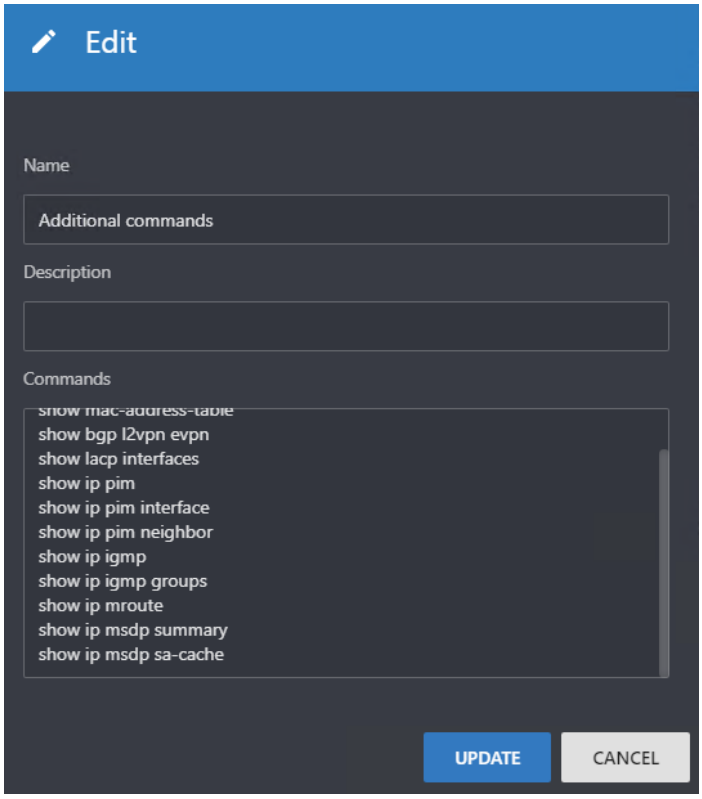

Enable and create the additional commands in NetEdit as needed [Settings -> Validation -> Change Validation -> Command Mapping]

3) CONFIGURE VSX AND FABRIC INFRASTRUCTURE (OSPF, BGP)

NetEdit Plan For Spine Switches

Create a plan for the spine switches in NetEdit [Devices -> select spines -> Action -> Edit Running Config]

Figure 8. Select devices for plan

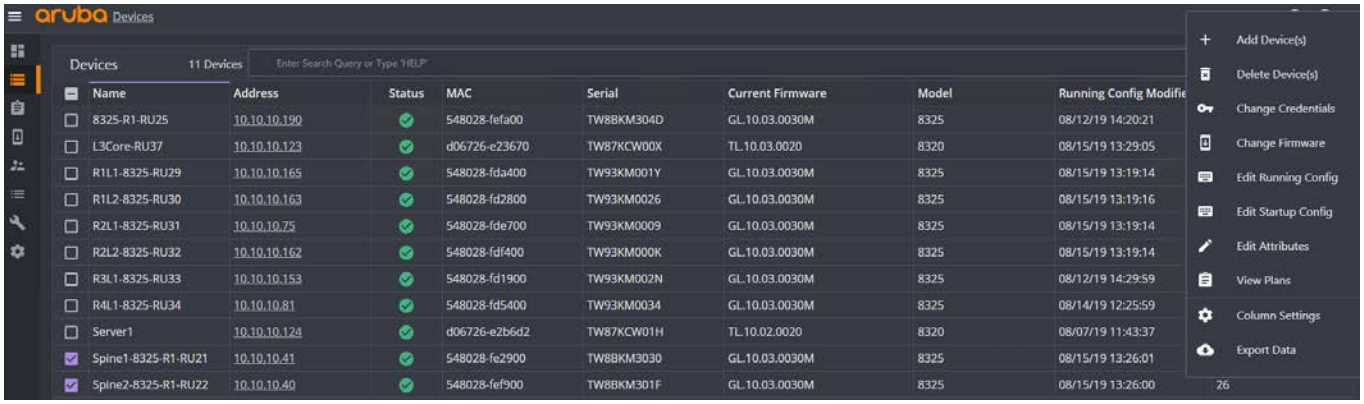

Give the plan a name and click “Create”.

Figure 9. Create Spine plan

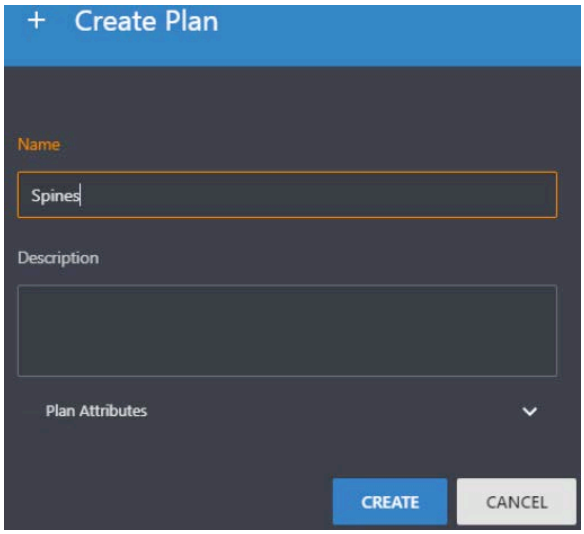

You should see the initial configs for your Spine switches in NetEdit.

The common configurations across the switches are shown as “white”, while the “blue” variables such as “HOSTNAME” and “A.B.C.D/M” have unique settings/values.

Figure 10. Common base configuration

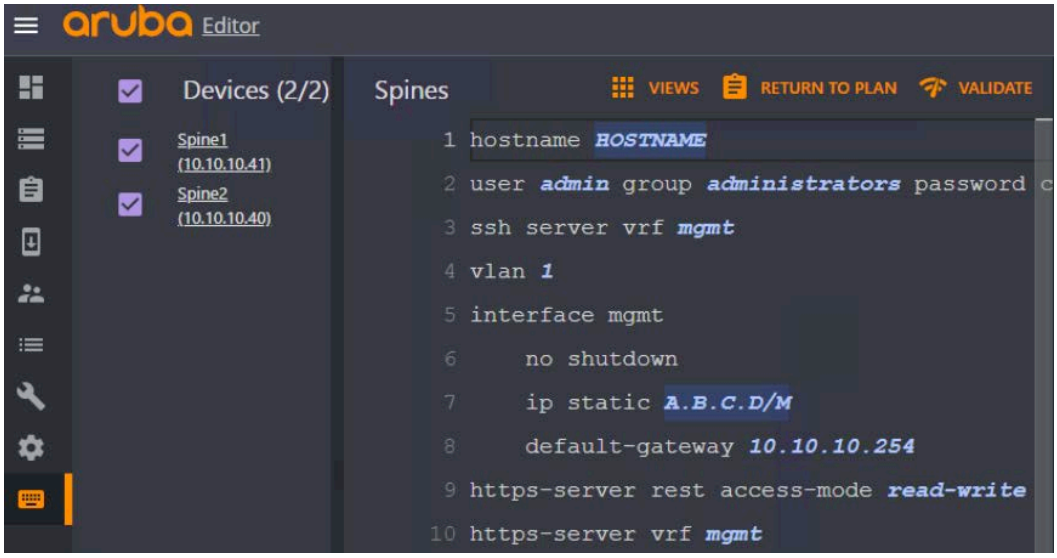

If you hover over “HOSTNAME” or “A.B.C.D/M” config (in blue), you can view the unique settings assigned to each switch.

Figure 11. Variables used for Spine mgmt

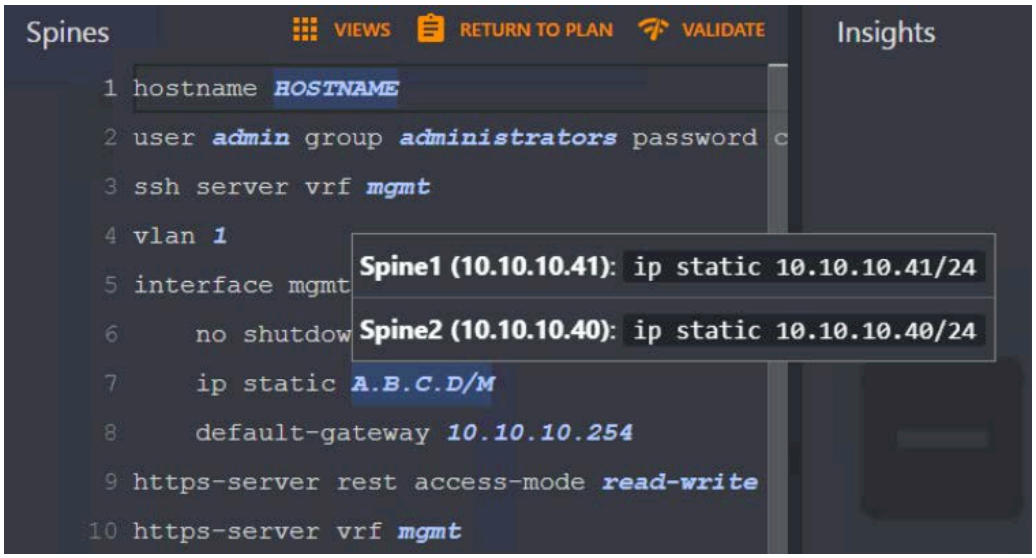

If you right click “HOSTNAME” or “A.B.C.D/M” (in blue), you will be able to modify the settings when in editing view.

Figure 12. Modify variables

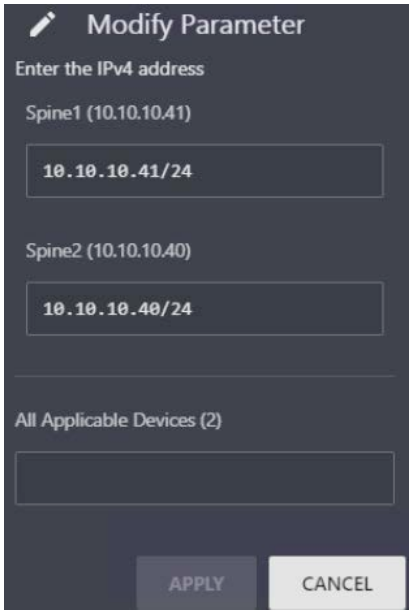

To begin the fabric infrastructure plan add the following OSPF and loopback configurations and right click to modify the loopbacks and router-IDs assigned to each spine.

```
router ospf 1
    router-id 1.1.1.1
    area 0.0.0.0
interface loopback 0
    ip address 1.1.1.1/32
    ip ospf 1 area 0.0.0.0
```

Figure 13. Adding OSPF/Loopbacks to Spines

```
4  router ospf 1
5      router-id 1.1.1.1
6      area 0.0.0.0
7  vlan 1
8  interface mgmt
9      no shutdown
10     ip static A.B.C.D/M
11     default-gateway 10.10.10.254
12 interface loopback 0
13     ip address 1.1.1.1/32
14     ip ospf 1 area 0
15 https-server rest access-mode read-write
16 https-server vrf mgmt
```

Figure 14. Modifying router-IDs

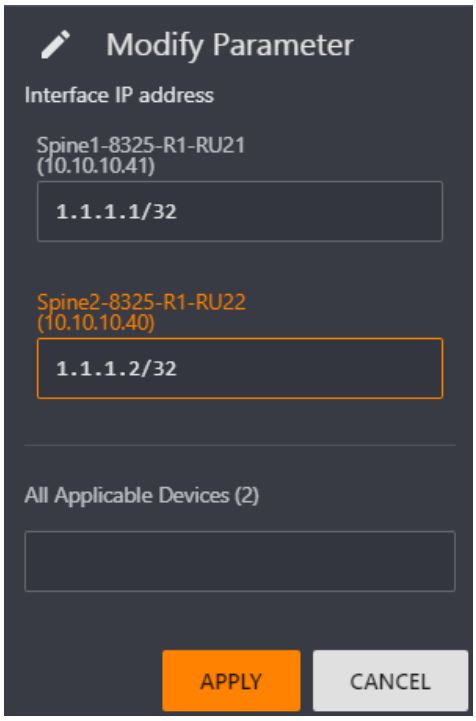

Figure 15. Modifying loopbacks

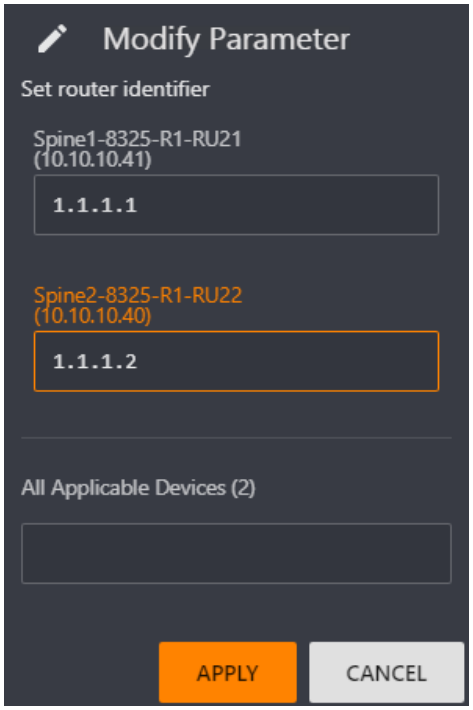

Next, configure the downlinks towards leaf switches and right click to modify the IPs for each downlink (not shown).

```
interface 1/1/1
    no shutdown
    ip address 101.1.1.0/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/2
    no shutdown
    ip address 101.1.1.4/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/3
    no shutdown
    ip address 101.1.1.8/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/4
    no shutdown
    ip address 101.1.1.10/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
```

Figure 16. Adding Fabric Interfaces

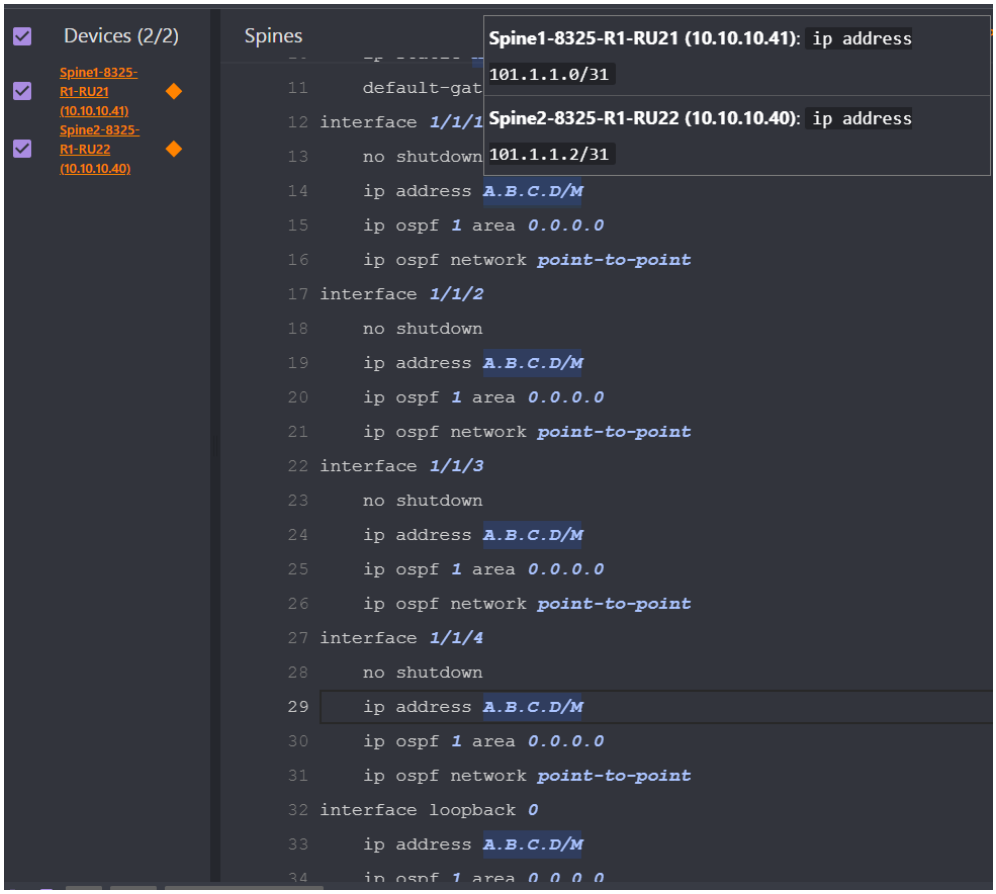

Next, add IBGP peering configurations towards all leaf switches and right click to modify the router-ID for each switch. The spine switches will act as Route Reflectors.

```
router bgp 65000
    bgp router-id 1.1.1.1
    neighbor 1.1.1.3 remote-as 65000
    neighbor 1.1.1.3 update-source loopback 0
    neighbor 1.1.1.4 remote-as 65000
    neighbor 1.1.1.4 update-source loopback 0
    neighbor 1.1.1.5 remote-as 65000
    neighbor 1.1.1.5 update-source loopback 0
    neighbor 1.1.1.6 remote-as 65000
    neighbor 1.1.1.6 update-source loopback 0
    address-family ipv4 unicast
        neighbor 1.1.1.3 activate
        neighbor 1.1.1.3 route-reflector-client
        neighbor 1.1.1.4 activate
        neighbor 1.1.1.4 route-reflector-client
```

```
neighbor 1.1.1.5 activate
neighbor 1.1.1.5 route-reflector-client
neighbor 1.1.1.6 activate
neighbor 1.1.1.6 route-reflector-client
```

Figure 17. Adding BGP

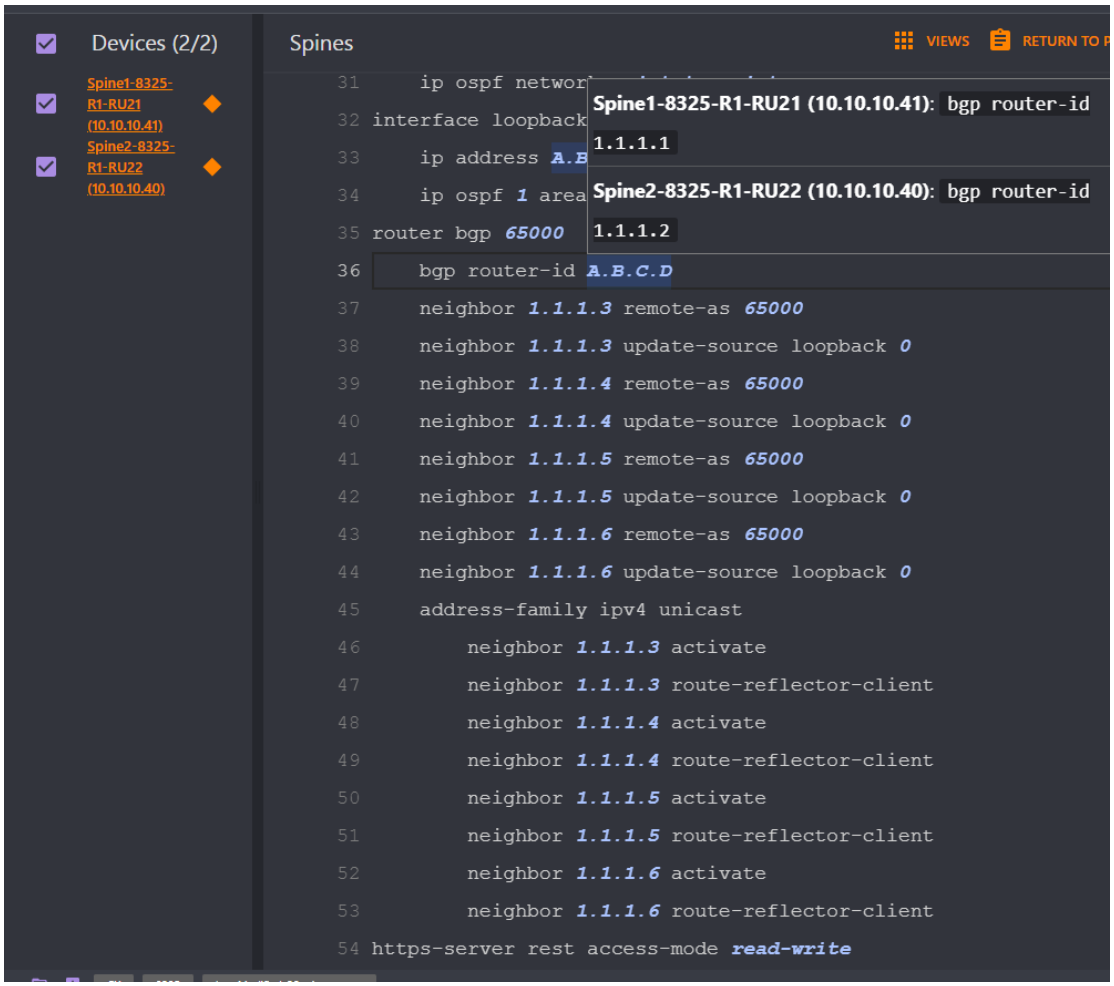

Finally, select “RETURN TO PLAN” -> “DEPLOY” -> “COMMIT” to push down and save the desired configs.

Figure 18. Deploy > Commit

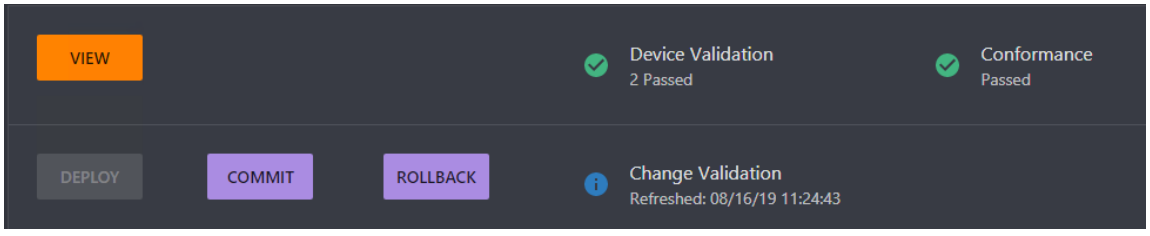

NetEdit Plan For VSX on Leaf Switches

Create a VSX plan for the leaf switches in NetEdit [Devices -> select leafs -> Action -> Edit Running Config] and name the plan.

Figure 19. Select VSX Leaf switches

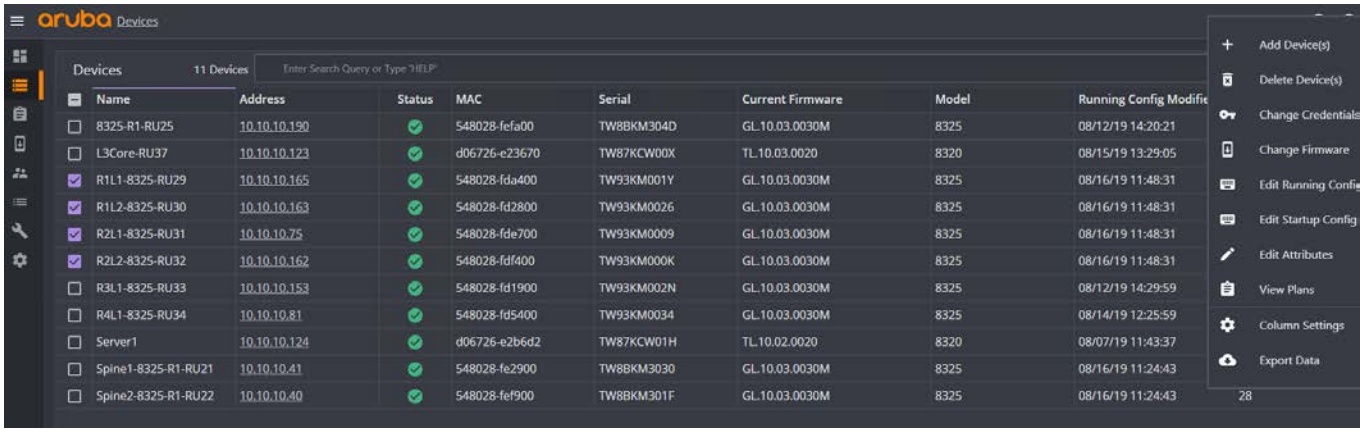

Configure VSX in both racks and right click to modify the interface IPs, keepalive settings, and roles for each device.

```
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
interface 1/1/47
    no shutdown
    ip address 111.1.1.1/24
interface 1/1/48
    no shutdown
    lag 1
vsx
    inter-switch-link lag 1
    role primary
    keepalive peer 111.1.1.2 source 111.1.1.1
    vsx-sync vsx-global
```

Figure 20. Adding VSX and modifying Keepalive settings

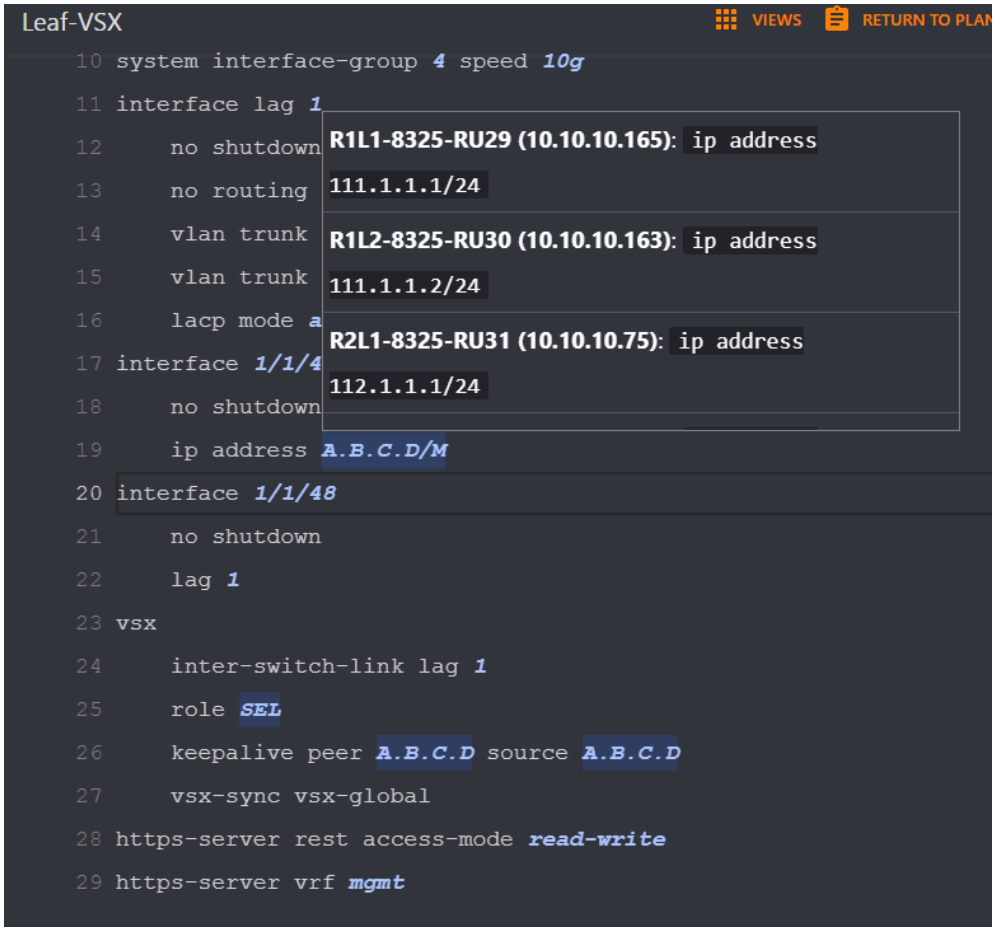

Select “RETURN TO PLAN” -> “DEPLOY” -> “COMMIT” to push down and save the desired configs.

Verify Change Validation

You can click on “Change Validation” to verify if VSX configuration worked as expected.

Figure 21. Rack 1 VSX status showing up and in-sync

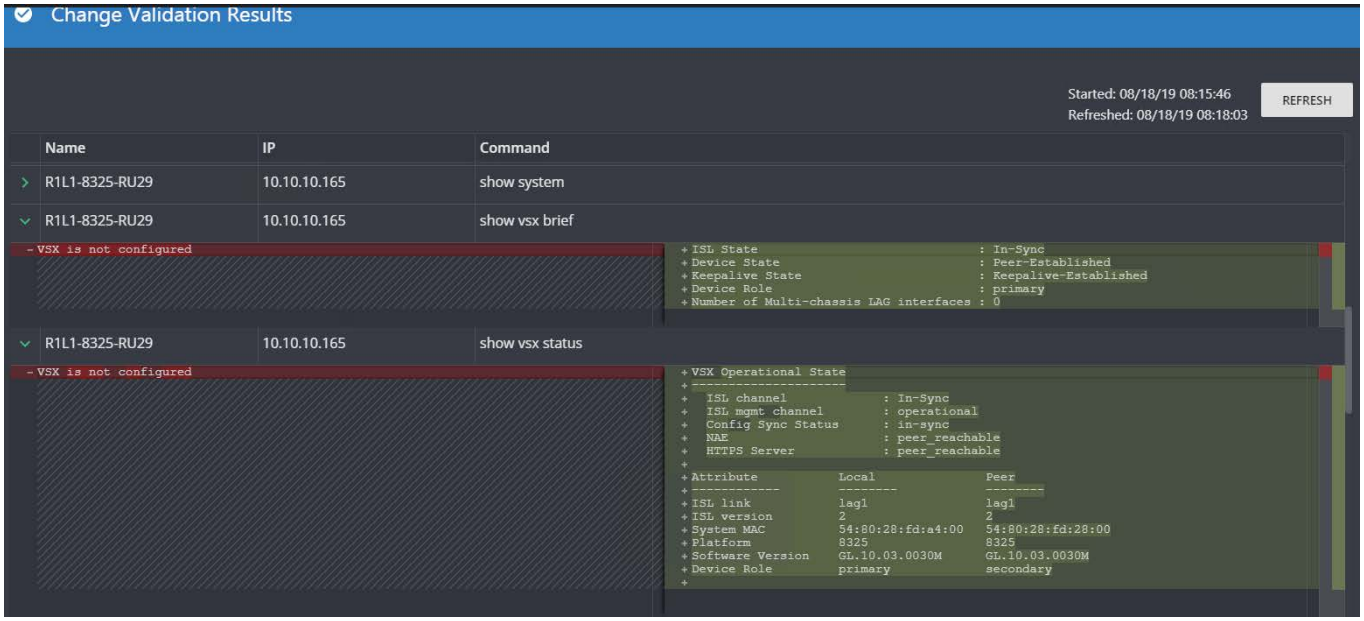

Figure 22. Rack 2 VSX status showing up and in-sync

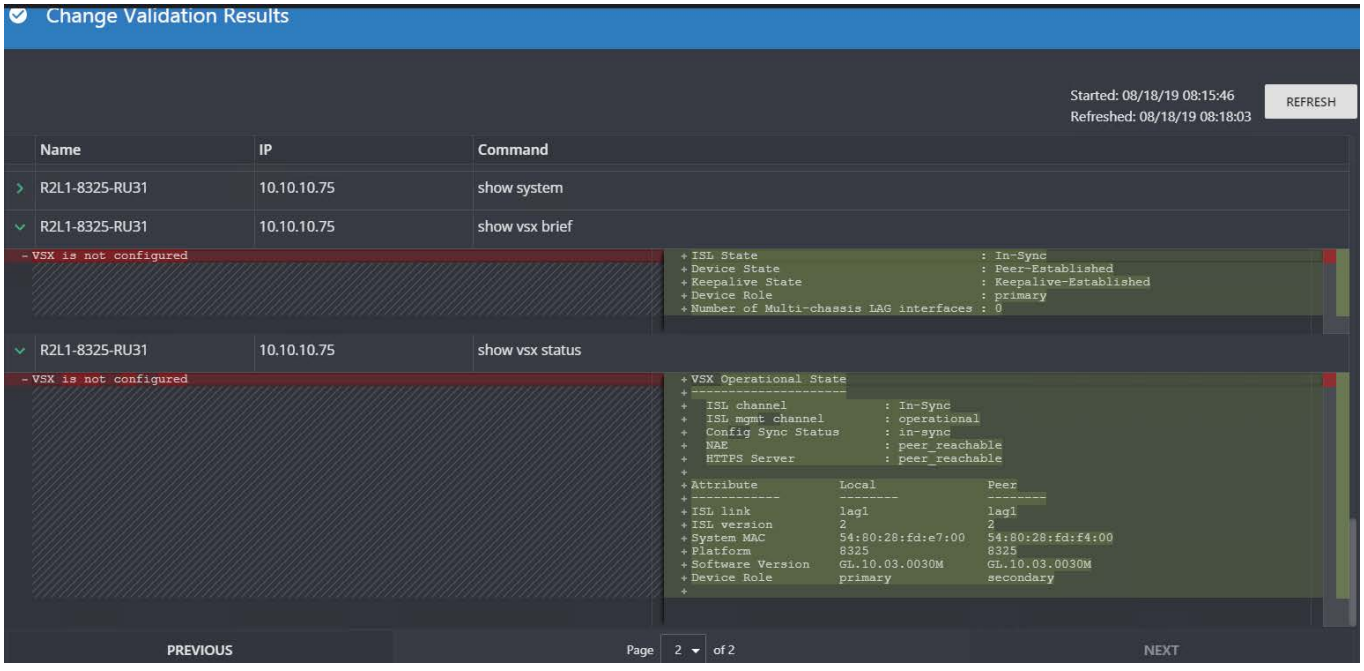

NetEdit Plan For Fabric Infrasructure on Leaf Switches

Create a Fabric Configuration plan for the leaf switches in NetEdit [Devices -> select the leafs -> Action -> Edit Running Config] and name the plan.

Figure 23. Selecting Leafs for Fabric Plan

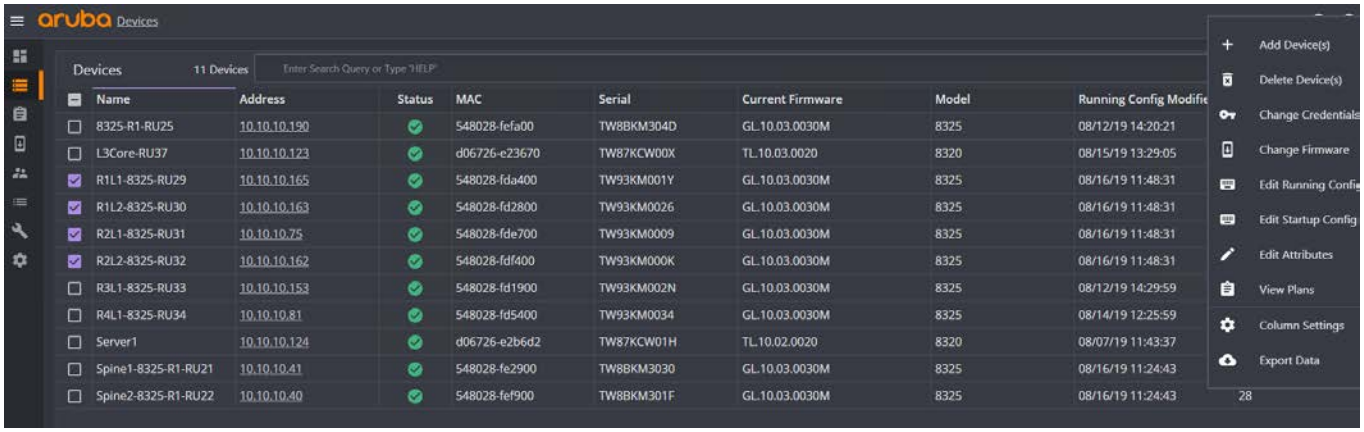

Start by adding the OSPF configuration and loopbacks. Right click to modify loopbacks and router-IDs assigned to each leaf.

```
router ospf 1
    router-id 1.1.1.3
    area 0.0.0.0
interface loopback 0
    ip address 1.1.1.3/32
    ip ospf 1 area 0.0.0.0
```

Figure 24. Adding OSPF and modifying router-IDs

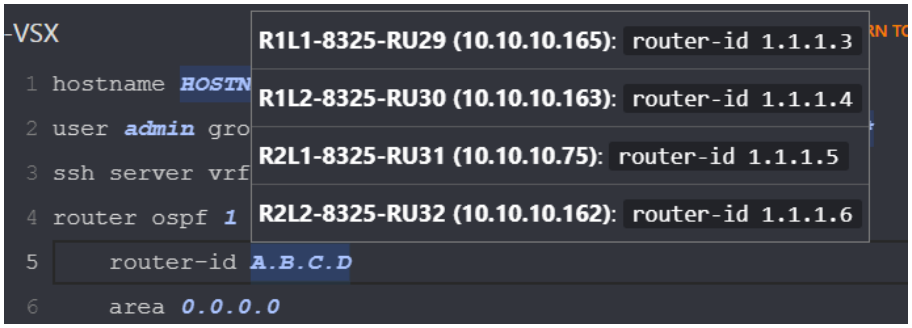

Figure 25. Adding and modifying loopbacks

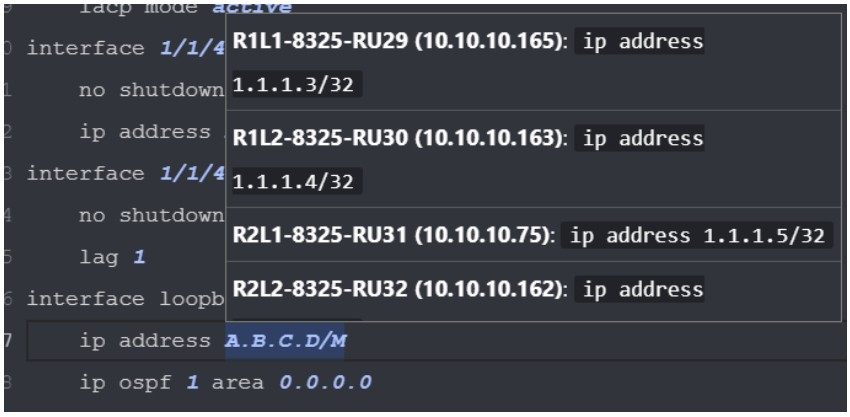

Next, configure the uplinks towards spine switches and right click to modify the IPs for each uplink.

```
interface 1/1/49
    no shutdown
    ip address 101.1.1.1/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/50
    no shutdown
    ip address 101.1.1.3/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
```

Figure 26. Adding and modifying uplink configurations

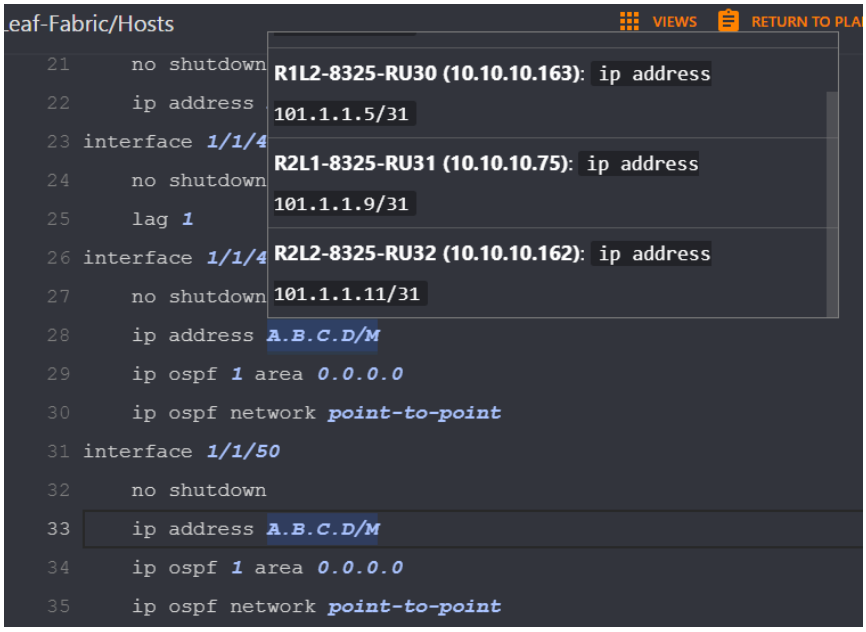

Now add the IBGP peering configurations towards all spine switches that function. Right click to modify the router-ids.

```
router bgp 65000
    bgp router-id 1.1.1.3
    neighbor 1.1.1.1 remote-as 65000
    neighbor 1.1.1.1 update-source loopback 0
    neighbor 1.1.1.2 remote-as 65000
    neighbor 1.1.1.2 update-source loopback 0
    address-family ipv4 unicast
        neighbor 1.1.1.1 activate
        neighbor 1.1.1.2 activate
```

Figure 27. Adding BGP and modifying router-IDs

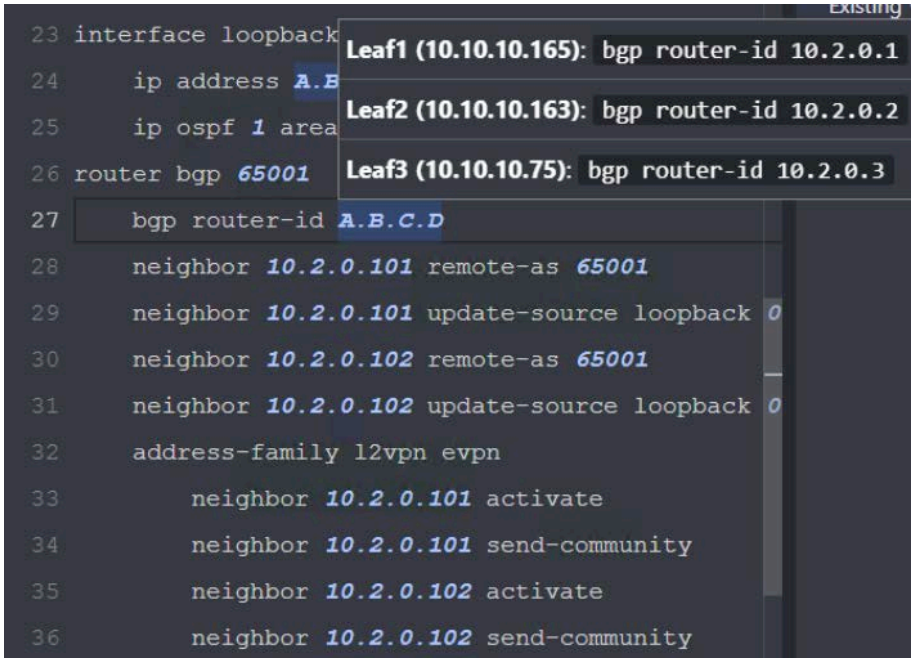

Finally, select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to verify if BGP and OSPF peering is working as expected.

Figure 28. Rack 2 Leaf 2 showing BGP in established state

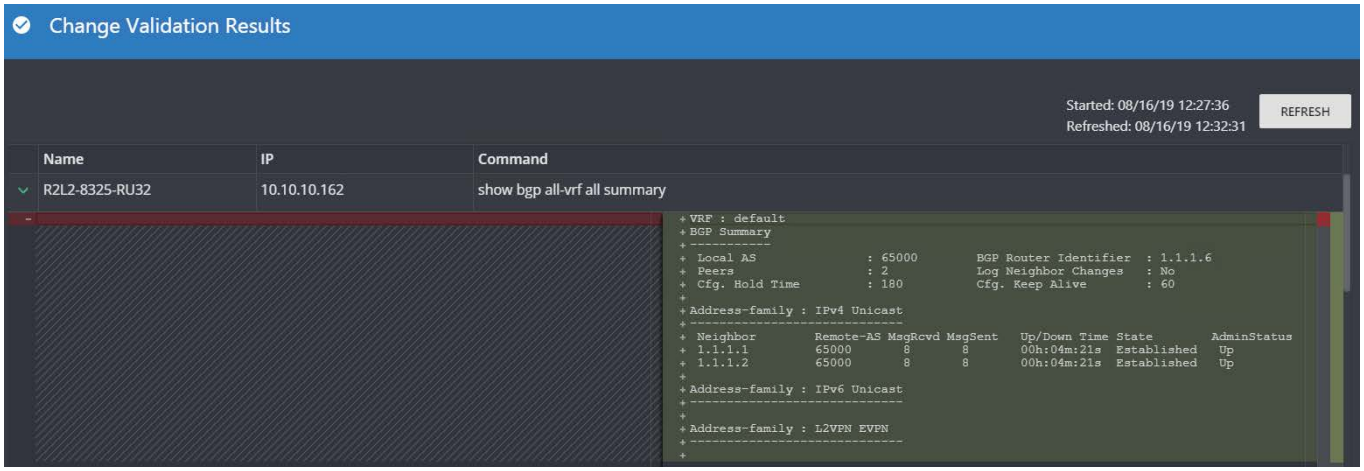

Figure 29. Rack 2 Leaf 2 showing OSPF in Full state

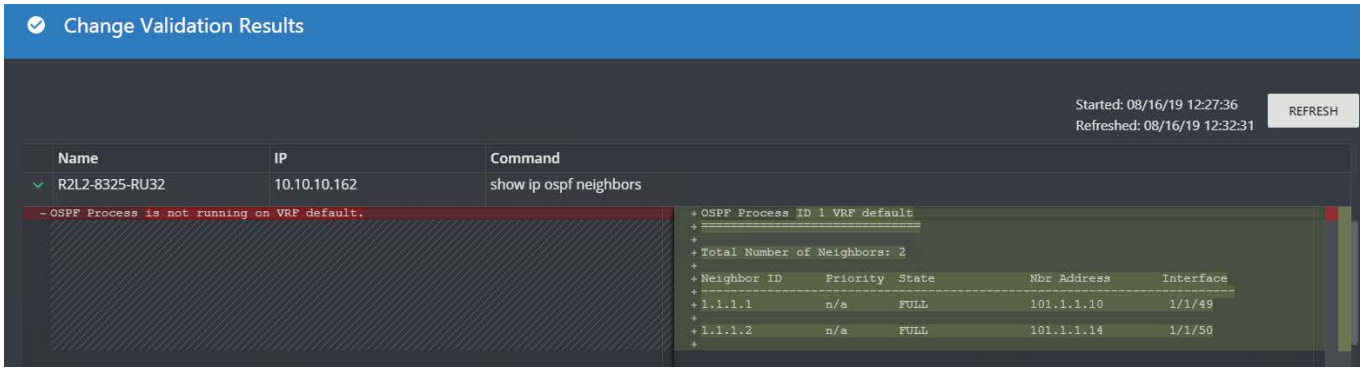

Figure 30. Rack 1 Leaf 2 showing BGP in established state

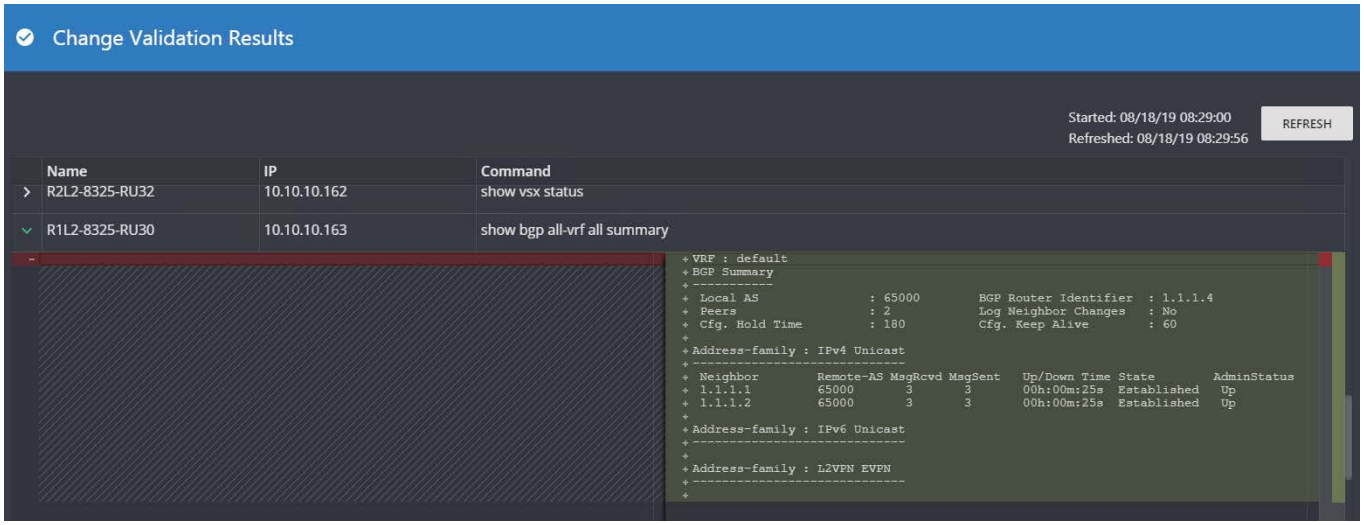

Figure 31. Rack 1 Leaf 2 showing OSPF in Full state

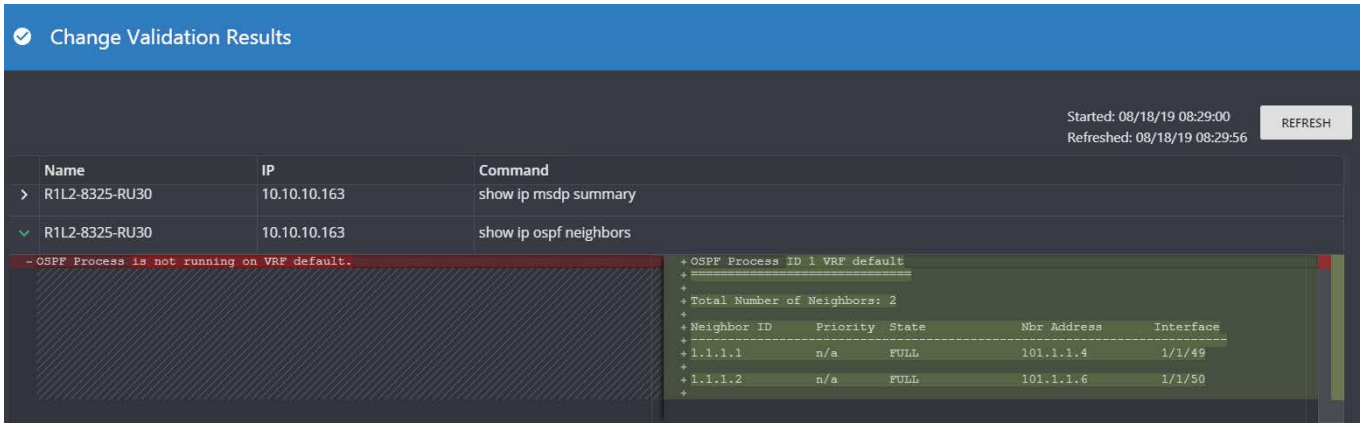

Figure 32. OSPF and BGP status on Spine 1

```
Spine1-8325-R1-RU21# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----------
 Local AS               : 65000        BGP Router Identifier  : 1.1.1.1
 Peers                  : 4            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State       AdminStatus
 1.1.1.3         65000       33      32      00h:26m:18s  Established   Up
 1.1.1.4         65000       33      33      00h:26m:18s  Established   Up
 1.1.1.5         65000       32      33      00h:26m:18s  Established   Up
 1.1.1.6         65000       34      33      00h:26m:18s  Established   Up

Spine1-8325-R1-RU21# sh ip ospf neighbors
OSPF Process ID 1 VRF default
================================

Total Number of Neighbors: 4

Neighbor ID      Priority  State           Nbr Address       Interface
-------------------------------------------------------------------------
1.1.1.3          n/a       FULL            101.1.1.1          1/1/1

1.1.1.4          n/a       FULL            101.1.1.5          1/1/2

1.1.1.5          n/a       FULL            101.1.1.9          1/1/3

1.1.1.6          n/a       FULL            101.1.1.11         1/1/4
```

After validation, you can choose to “COMMIT” to save the desired configs or “ROLLBACK” to revert configs before the configs were deployed and make further desired changes.

4) CONFIGURE HOST INFRASTRUCTURE (VLANS AND LAGS)

NetEdit Plan For Host Configuration on Leaf Switches

Create a Host Configuration plan for the leaf switches in NetEdit [Devices -> select all leafs -> Action -> Edit Running Config] and name the plan.

Give the plan a name and “Create”.

Figure 33. Selecting Leafs for Host Plan

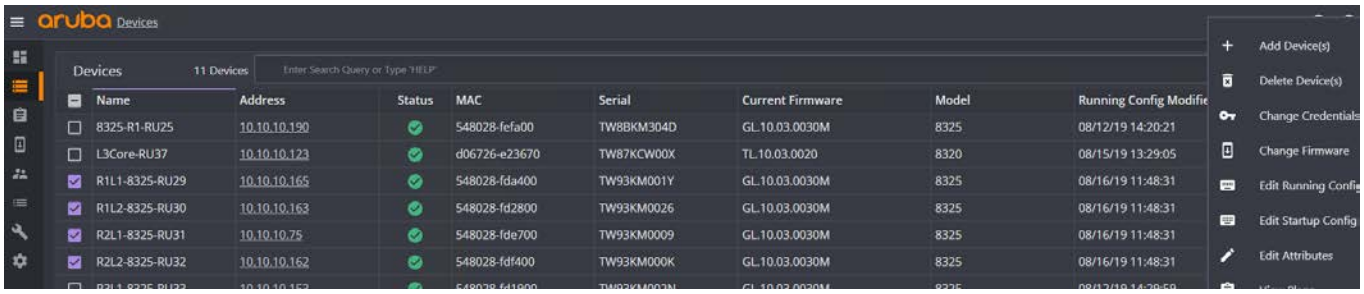

Start the configuration by choosing the appropriate Leaf switches from the left panel and then applying the required VLANs and VLAN Interfaces. Right click and modify the VLAN IP addresses as necessary.

```
# Rack 1

vlan 10
interface vlan10
    vsx-sync active-gateways
    ip address 10.1.1.253/24
    active-gateway ip 10.1.1.254 mac 00:00:00:00:00:01
    ip ospf 1 area 0.0.0.0
```

Figure 34. Adding VLAN 10 to Rack 1

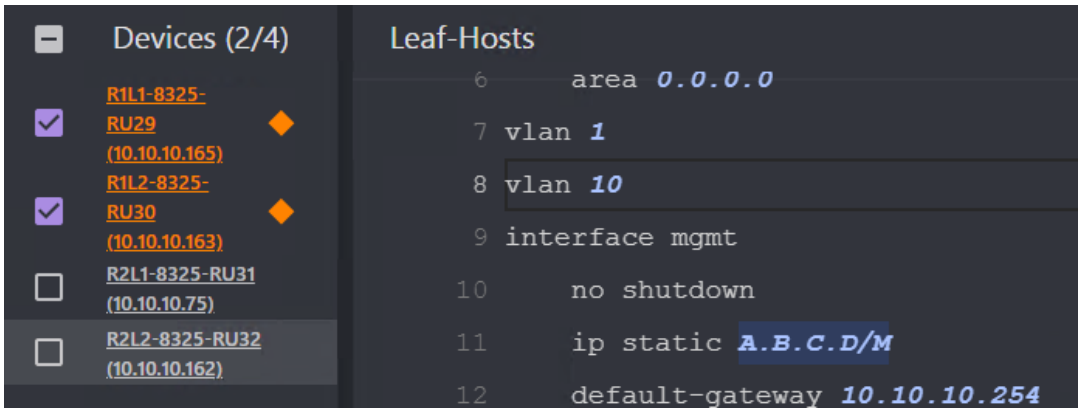

Figure 35. Adding VLAN 10 Interface to Rack 1

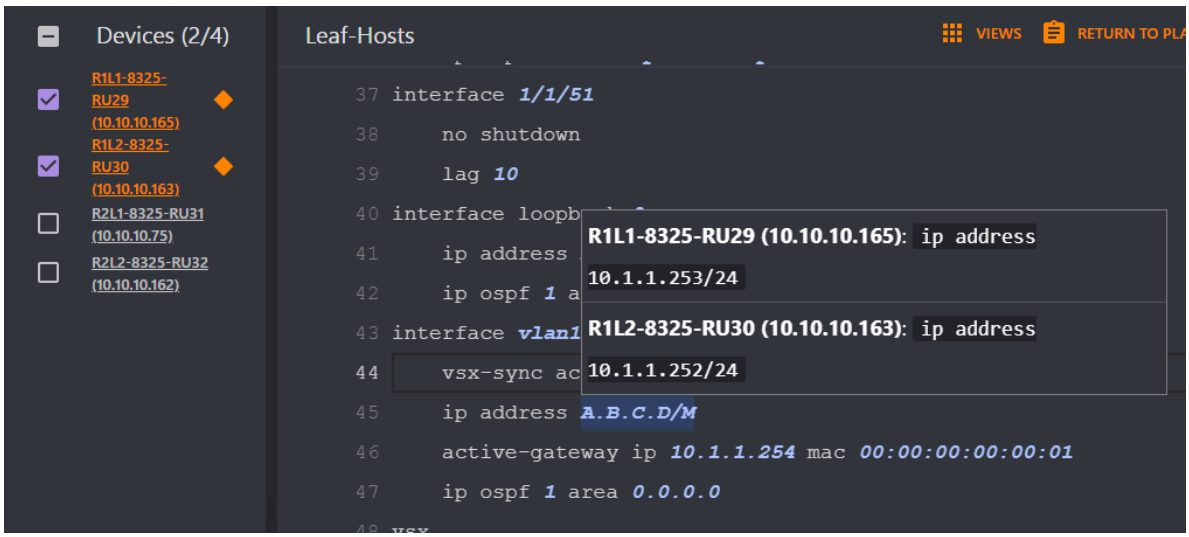


```
# Rack 2

vlan 11
interface vlan11
    vsx-sync active-gateways
    ip address 11.1.1.253/24
    active-gateway ip 11.1.1.254 mac 00:00:00:00:00:03
    ip ospf 1 area 0.0.0.0
```

Figure 36. Adding VLAN 11 to Rack 2

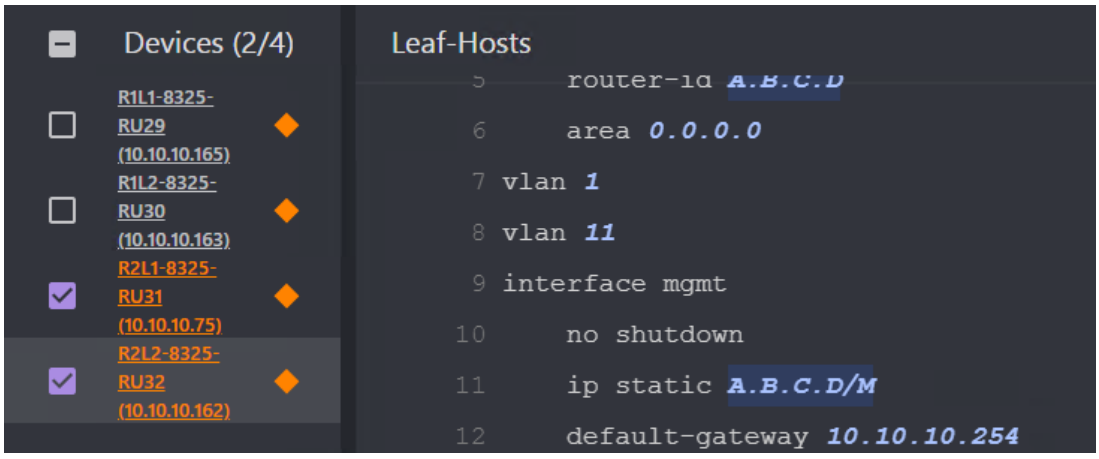

Figure 37. Adding VLAN 11 Interface to Rack 2

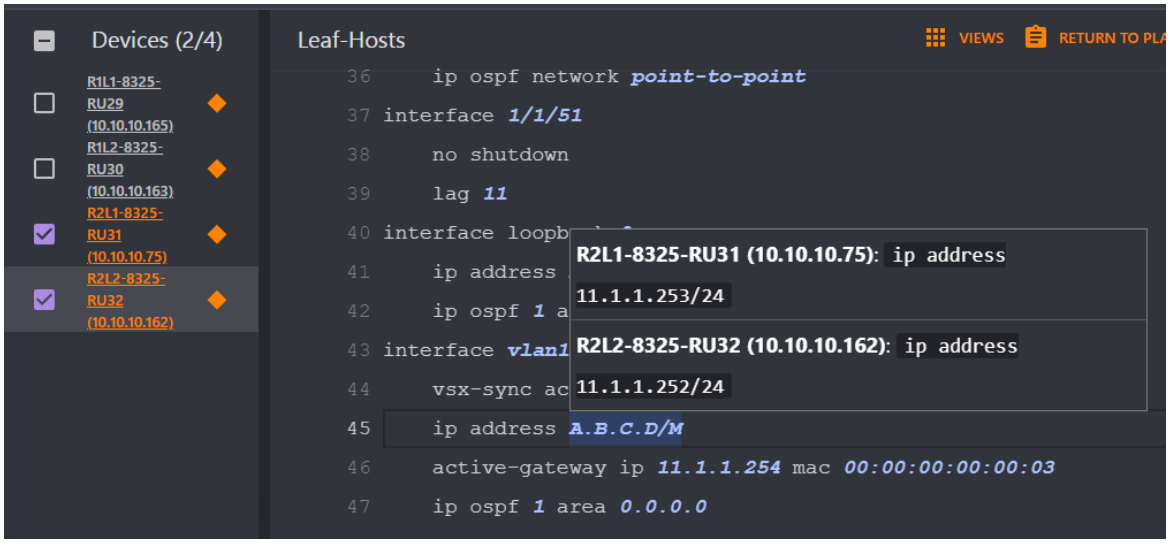

Next, select the appropriate leaf switches facing the servers and crate the needed LAG interfaces and apply to the appropriate interface.

```
# Rack 1

interface lag 10 multi-chassis
    vsx-sync vlans
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 10
    lacp mode active
interface 1/1/51
    no shutdown
    lag 10


# Rack 2

interface lag 11 multi-chassis
    vsx-sync vlans
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 11
    lacp mode active
interface 1/1/51
    no shutdown
```

```
lag 11
```

Figure 38. View after MLAGs were added to each device

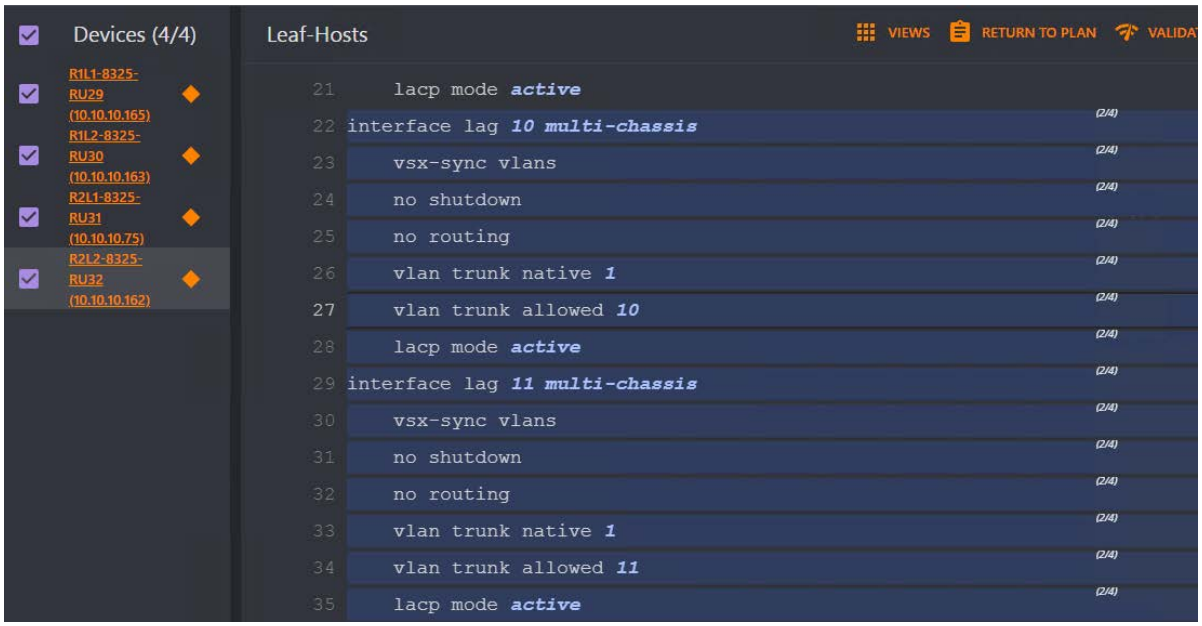

Figure 39. View of 1/1/51 showing MLAG assignments

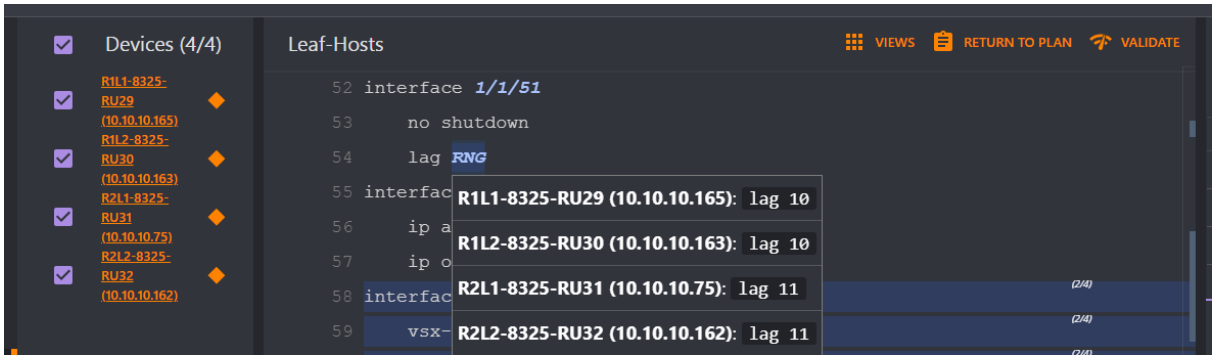

Finally, select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to verify if the VLANs, Interfaces, and LAGs are up.

Figure 40. Show interface brief view of Rack 2 Leaf 2 switch showing VLANs, and host interfaces are up

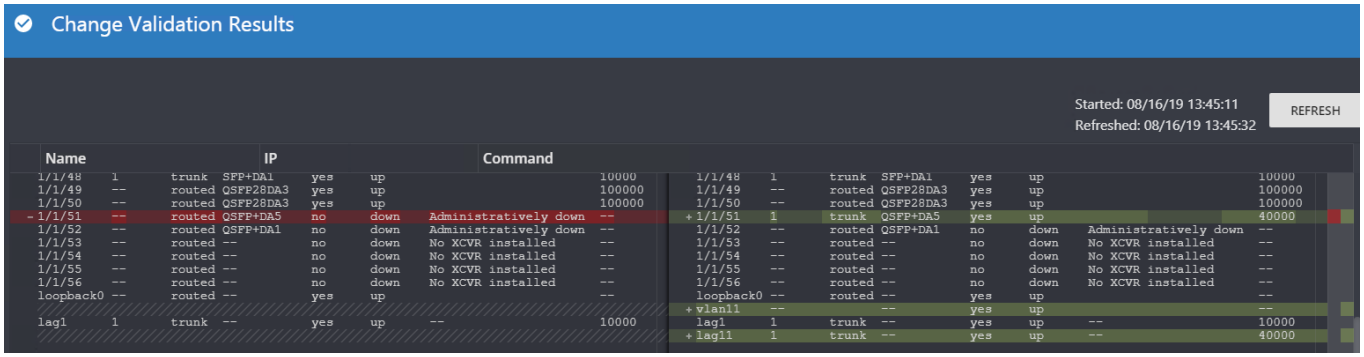

Figure 41. Show LACP interface view of Rack 2 Leaf 2 switch showing LACP is up and active

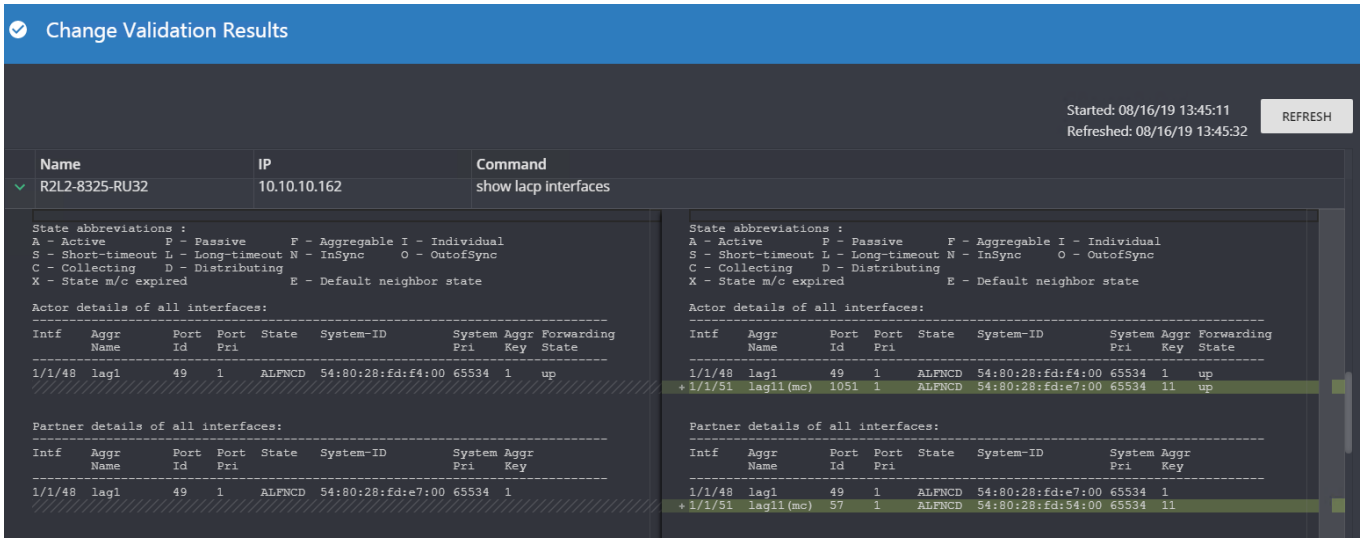

5) CONFIGURE BORDER CONNECTION TO L3CORE ENVIRONMENT

NetEdit Plan For Leaf Switch Border Connection in Rack 2 and L3Core

Create a Host Configuration plan for Rack 2 leaf switches and the L3Core in NetEdit [Devices -> select rack 2 leafs and L3 core-> Action -> Edit Running Config] and name the plan.

Give the plan a name and “Create”.

Figure 42. Selecting switches for Border to L3Core connection

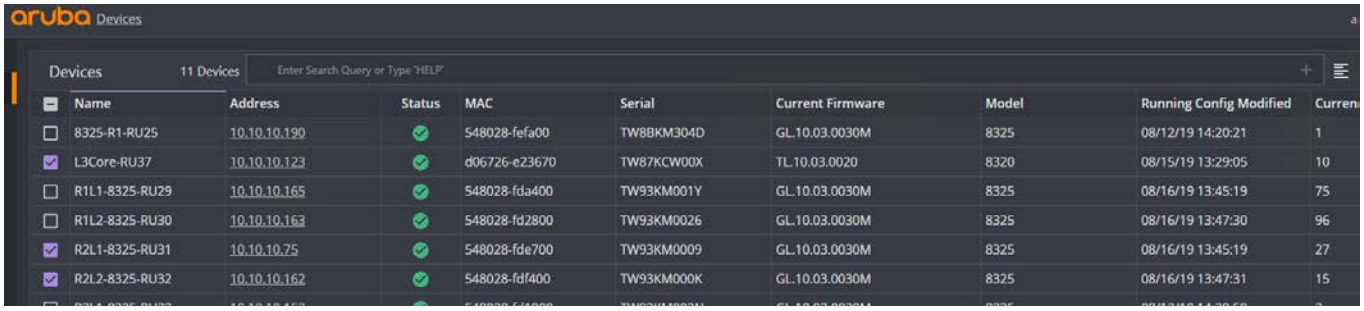

Name	Address	Status	MAC	Serial	Current Firmware	Model	Running Config Modified	Curren
8325-R1-RU25	10.10.10.190		548028-fefa00	TW8BKM304D	GL.10.03.0030M	8325	08/12/19 14:20:21	1
L3Core-RU37	10.10.10.123		d06726-e23670	TW87KCW00X	TL.10.03.0020	8320	08/15/19 13:29:05	10
R1L1-8325-RU29	10.10.10.165		548028-fda400	TW93KM001Y	GL.10.03.0030M	8325	08/16/19 13:45:19	75
R1L2-8325-RU30	10.10.10.163		548028-fd2800	TW93KM0026	GL.10.03.0030M	8325	08/16/19 13:47:30	96
R2L1-8325-RU31	10.10.10.75		548028-fde700	TW93KM0009	GL.10.03.0030M	8325	08/16/19 13:45:19	27
R2L2-8325-RU32	10.10.10.162		548028-fdf400	TW93KM000K	GL.10.03.0030M	8325	08/16/19 13:47:31	15

Choose the Rack 2 Leaf switches and configure the Physical interface and IP addresses. Right click and modify the IP addresses as necessary.

```
# Rack 2

interface 1/1/52
    no shutdown
    ip address 111.1.1.3/31
router bgp 65000
    neighbor 111.1.1.2 remote-as 65001
    address-family ipv4 unicast
        neighbor 111.1.1.2 activate
        network 10.1.1.0/24
        redistribute ospf
router ospf 1
    redistribute bgp
```

Figure 43. BGP peering to L3Core added to Rack 2

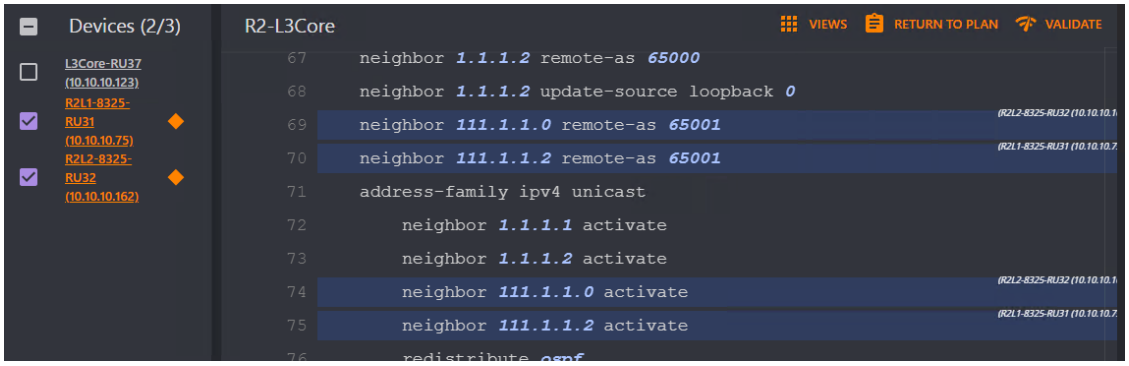

Figure 44. L3 Interfaces to L3Core added

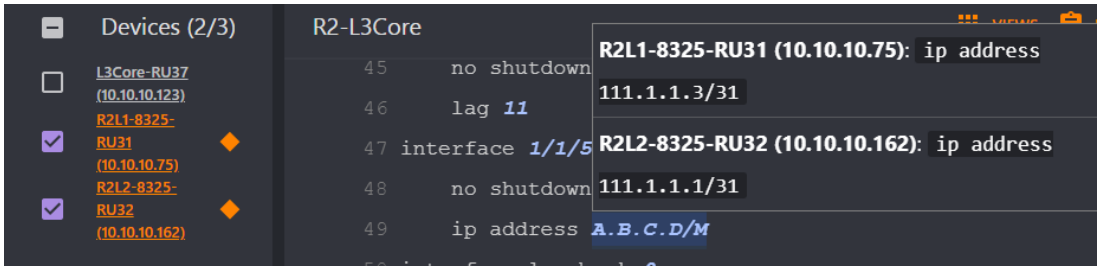

Figure 45. Redistribute command added between OSPF/BGP

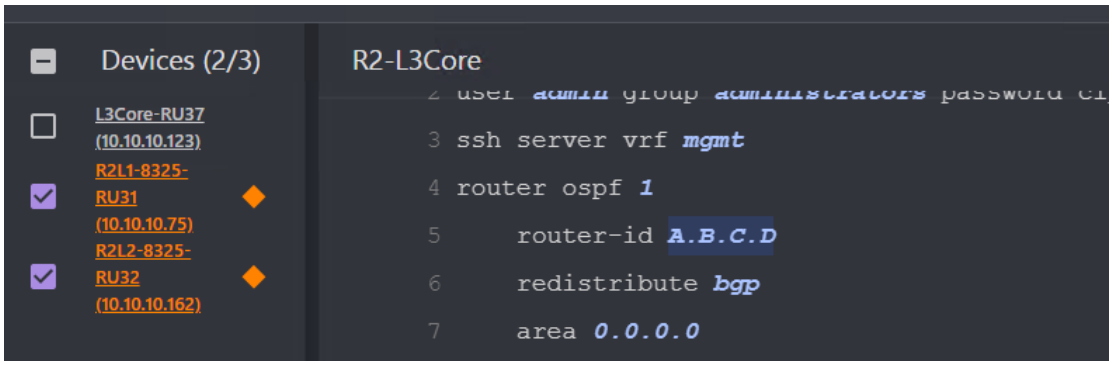

Choose the L3Core and configure interfaces and BGP peering.

```
# L3Core

interface 1/1/49
    no shutdown
    ip address 111.1.1.0/31
interface 1/1/50
    no shutdown
    ip address 111.1.1.2/31
interface loopback 0
    ip address 1.1.1.7/32
router bgp 65001
    bgp router-id 1.1.1.7
    neighbor 111.1.1.1 remote-as 65000
    neighbor 111.1.1.3 remote-as 65000
    address-family ipv4 unicast
        neighbor 111.1.1.1 activate
        neighbor 111.1.1.3 activate
        network 1.1.1.7/32
```

Figure 46. BGP and Fabric interfaces added to L3Core

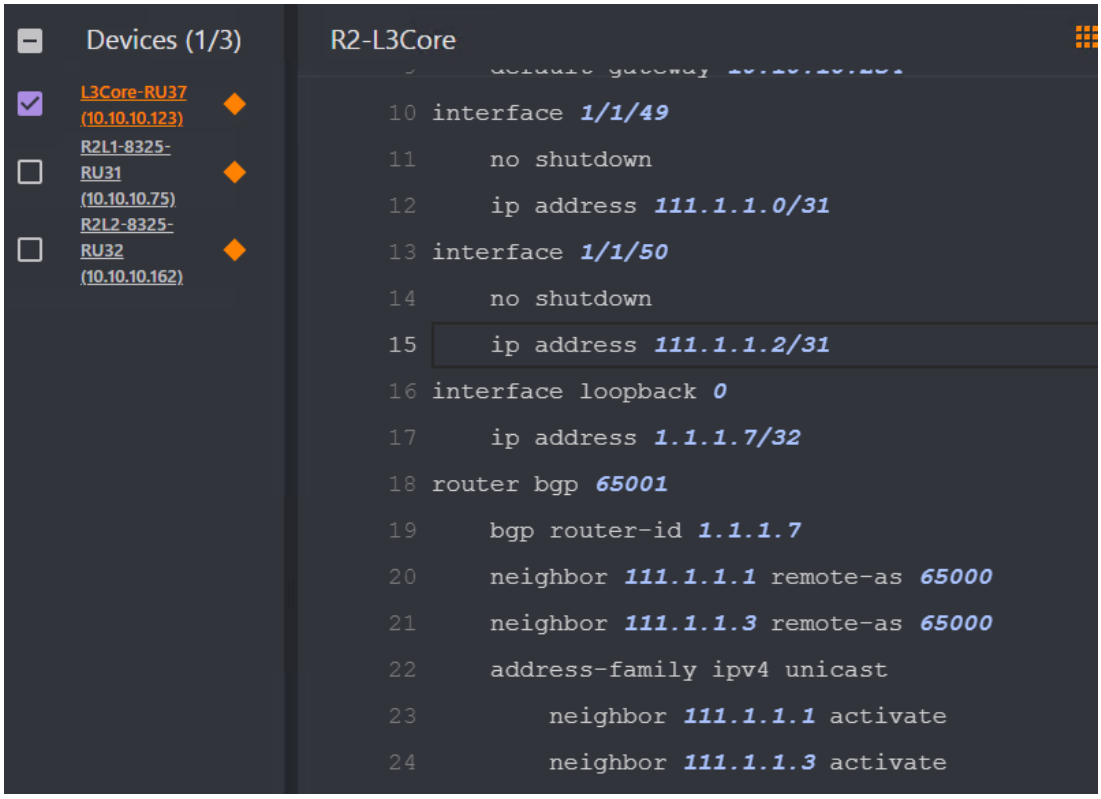

Then choose the L3Core and configure the host facing VLANs, addresses, and apply to BGP.

```
# L3Core

vlan 31
interface 1/1/1
    no shutdown
    speed 1000-full
    no routing
    vlan access 31
interface vlan31
    ip address 31.1.1.254/24
router bgp 65001
    address-family ipv4 unicast
    network 31.1.1.0/24
```

Figure 47. Host facing interface and VLAN added to L3Core

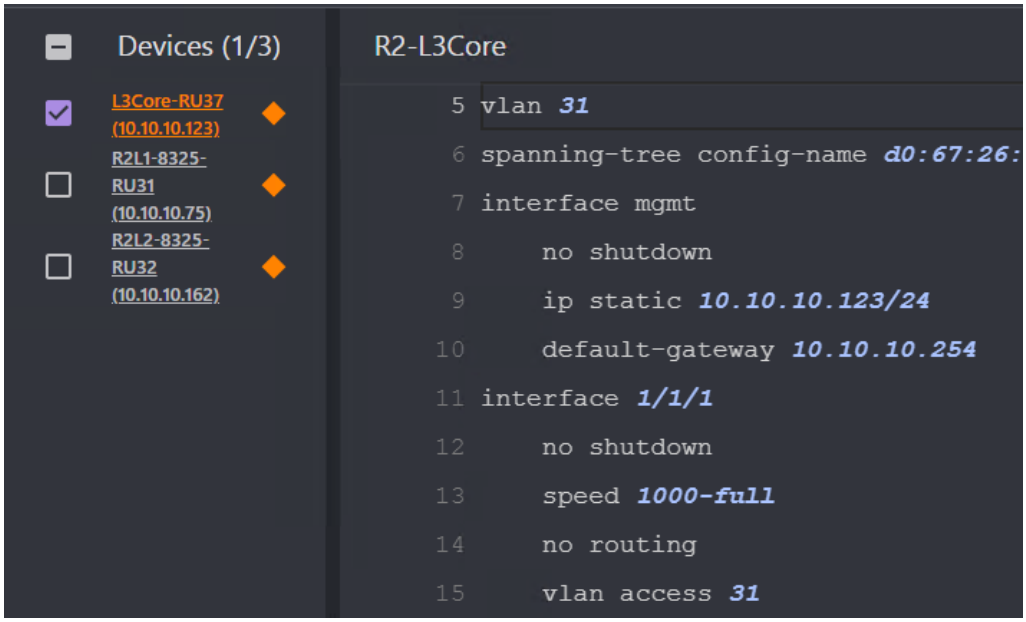

Figure 48. Adding VLAN interface and BGP configuration

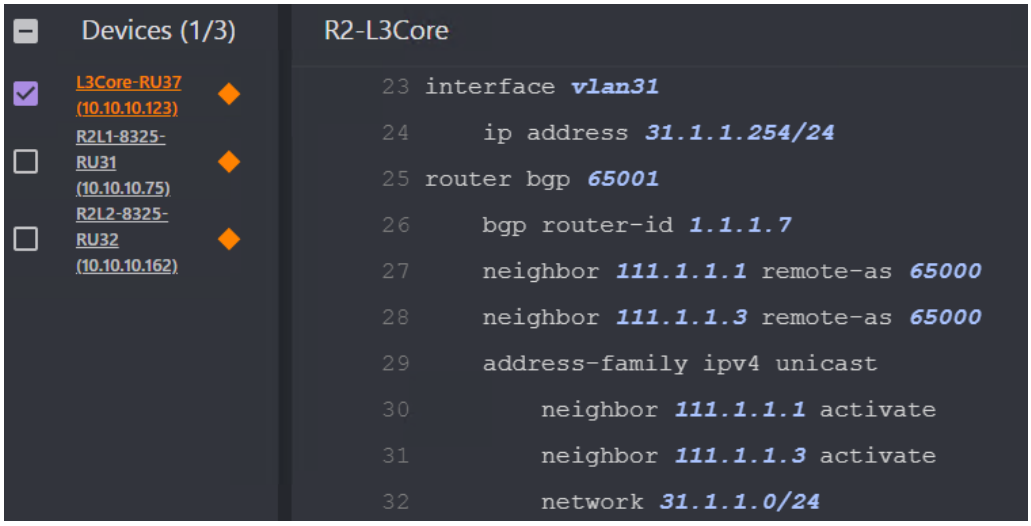

Finally, select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to verify if the VLANs, Interfaces, peering is up, and all routes can be seen.

Figure 49. L3Core showing Established BGP state with Rack 2

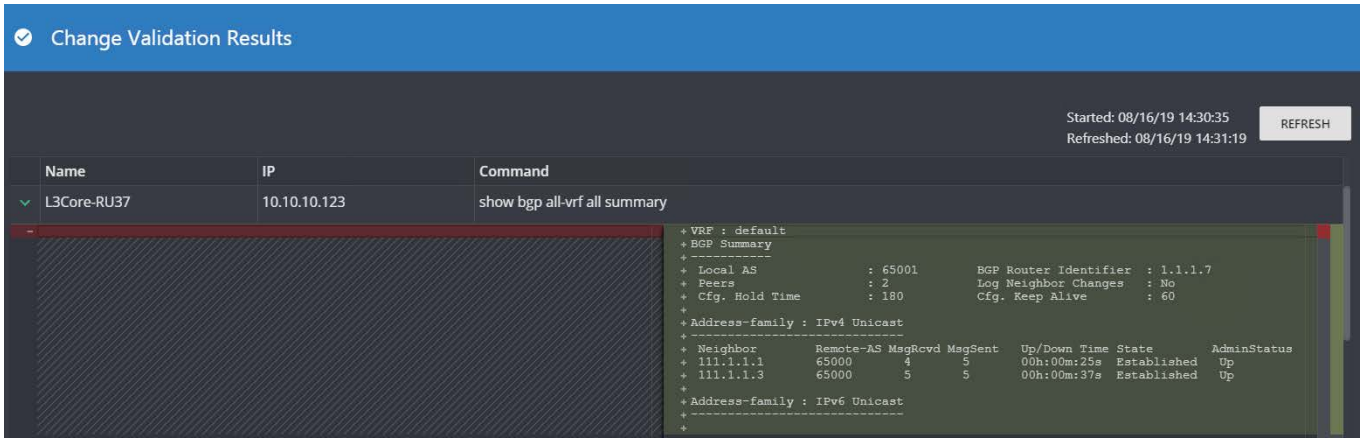

Figure 50. L3Core showing host facing interface is up and set as access for VLAN 31

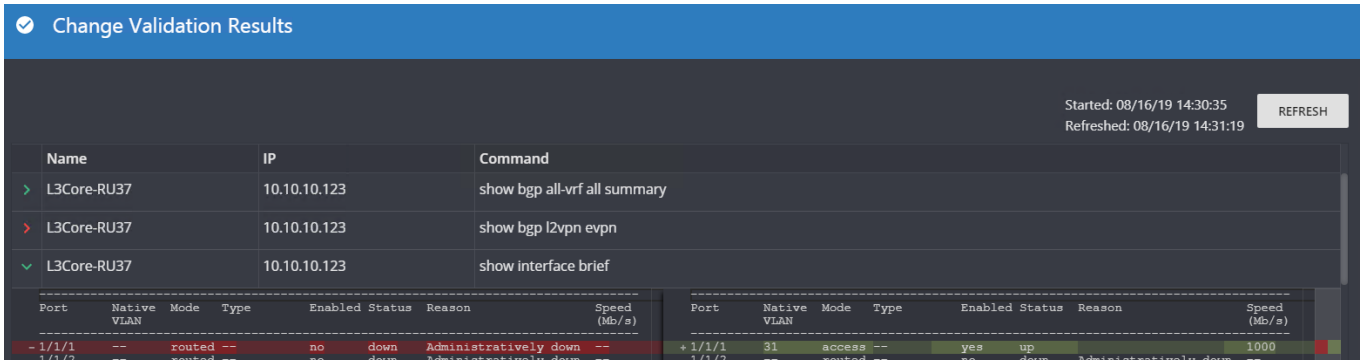

Figure 51. L3Core showing Rack 2 interfaces, Loopback, and VLAN 31 are up

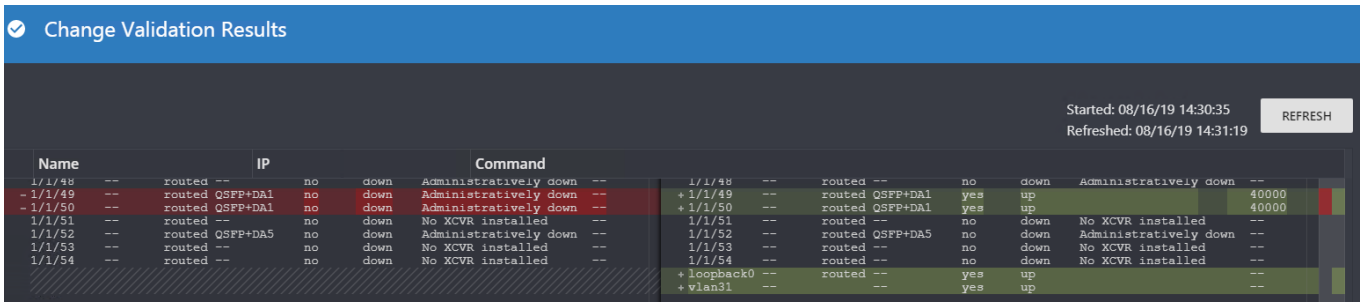

Figure 52. L3Core showing that it sees routes to VLANs and Loopbacks within the Spine/Leaf fabric

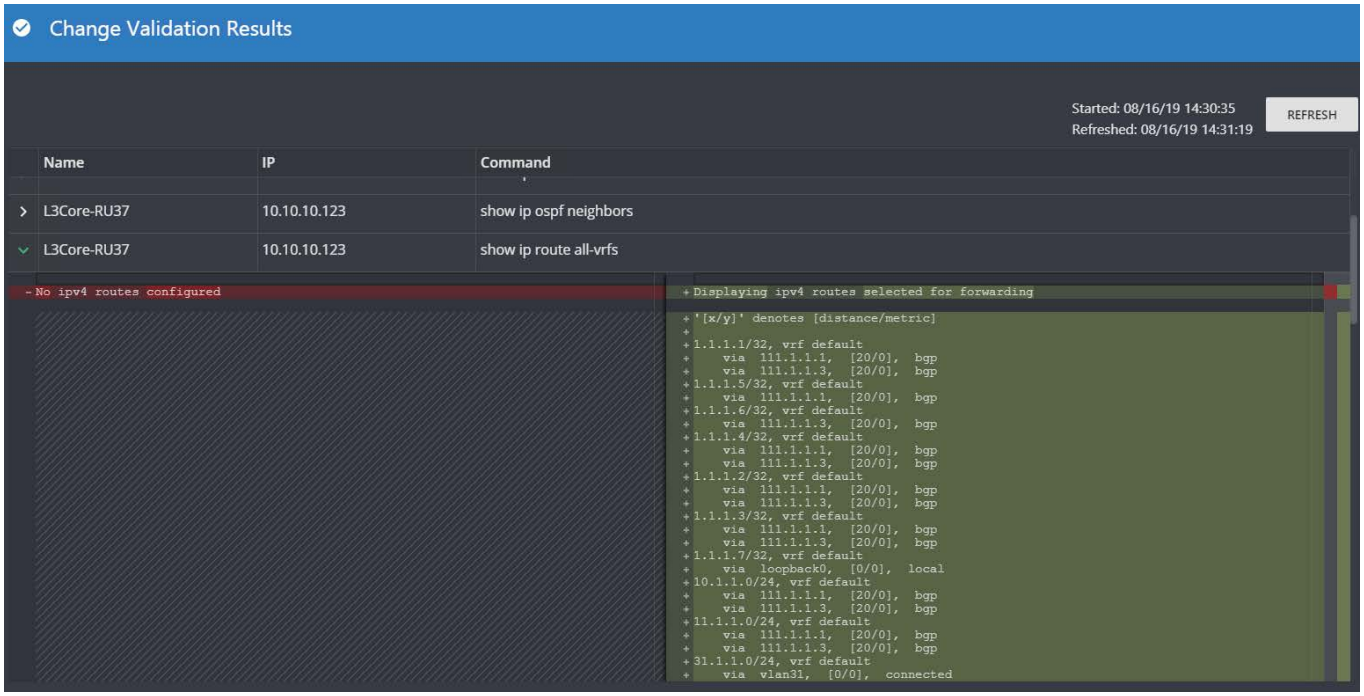

6) CONFIGURE MULTICAST ENVIRONMENT

NetEdit Plan For Multicast Configuration on Leaf Switches

Create a Multicast Leaf Configuration plan for Rack 1 & 2 leaf switches in NetEdit [Devices -> select rack 1 & 2 leafs -> Action -> Edit Running Config] and name the plan.

Give the plan a name and “Create”.

Figure 53. Selecting leaf switches for Multicast Plan

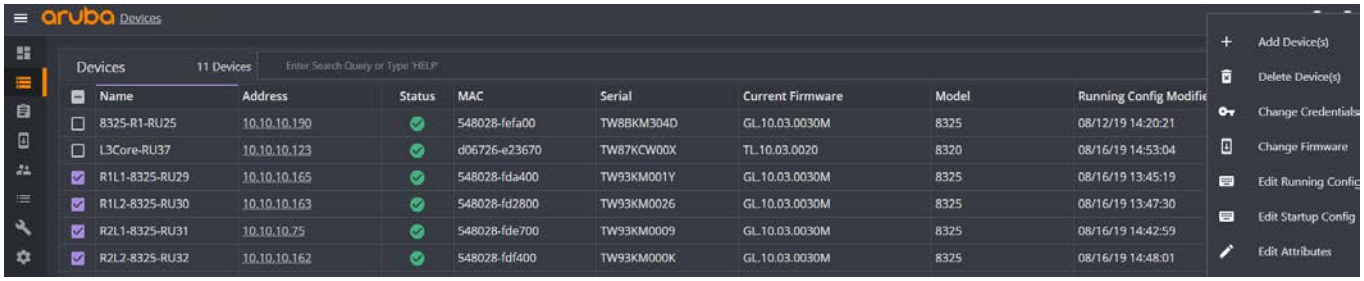

Start by choosing the appropriate Leaf switches from the left panel and apply the required VLANs and VLAN Interfaces. Right click and modify the VLAN IP address as necessary.

```
# All Leaf Switches
```

```
router pim
    enable
    rp-address 1.1.1.1
    active-active
interface 1/1/49
    ip pim-sparse enable
interface 1/1/50
    ip pim-sparse enable
interface loopback 0
    ip pim-sparse enable


# Rack 1 Leaf Switches


vlan 10
    ip igmp snooping enable
interface vlan10
    ip igmp enable
    ip pim-sparse enable


# Rack 2 Leaf Switches


interface 1/1/52
    ip pim-sparse enable
vlan 11
    ip igmp snooping enable
interface vlan11
    ip igmp enable
    ip pim-sparse enable
```

Figure 54. IGMP-snooping showing as added to VLANs 10 and 11

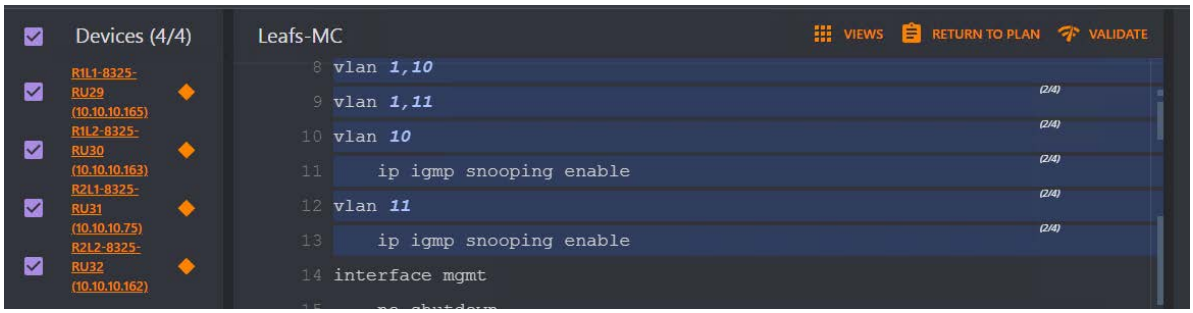

Figure 55. IP PIM-Sparse added to Fabric Interfaces and Loopback interfaces

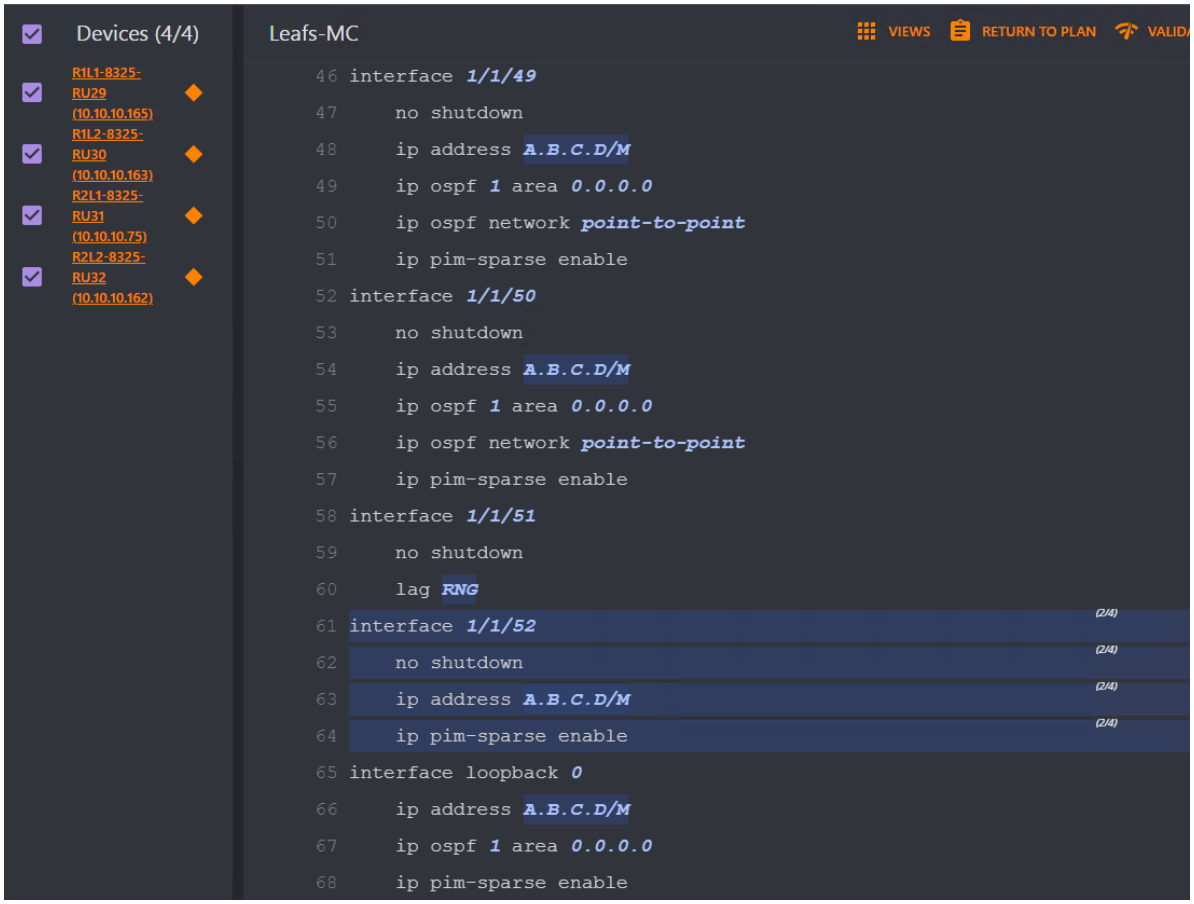

Figure 56. IP PIM-Sparse added to VLAN interfaces

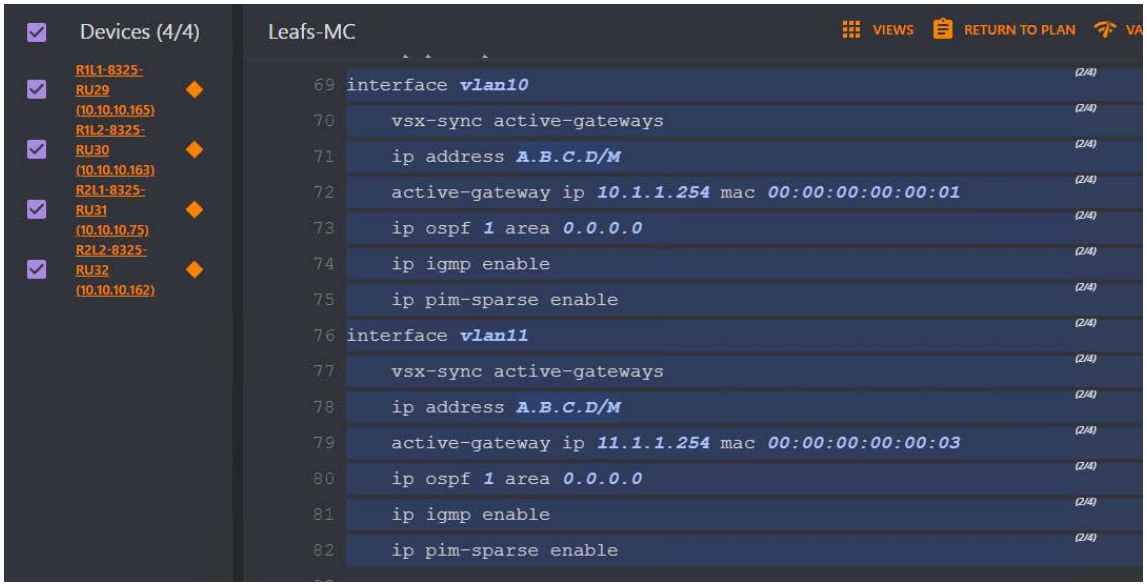

Figure 57. Enabling PIM globally and setting Spine 1 as the Static RP

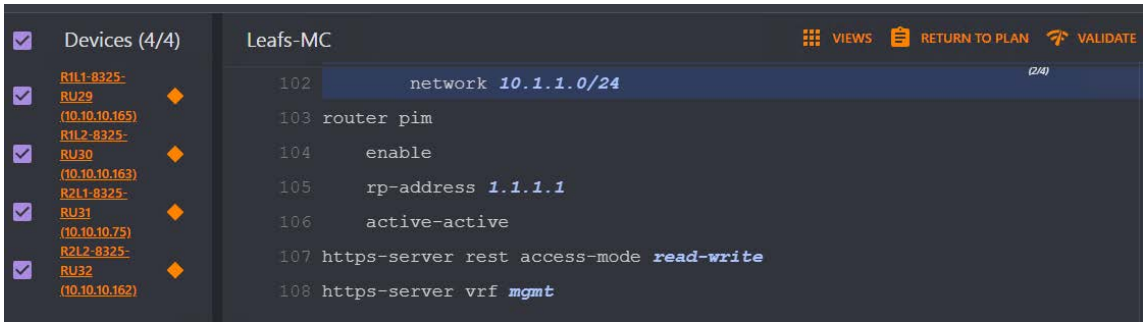

Finally, select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to verify if PIM and IGMP is up.

Figure 58. Rack2 showing IGMP is now configured

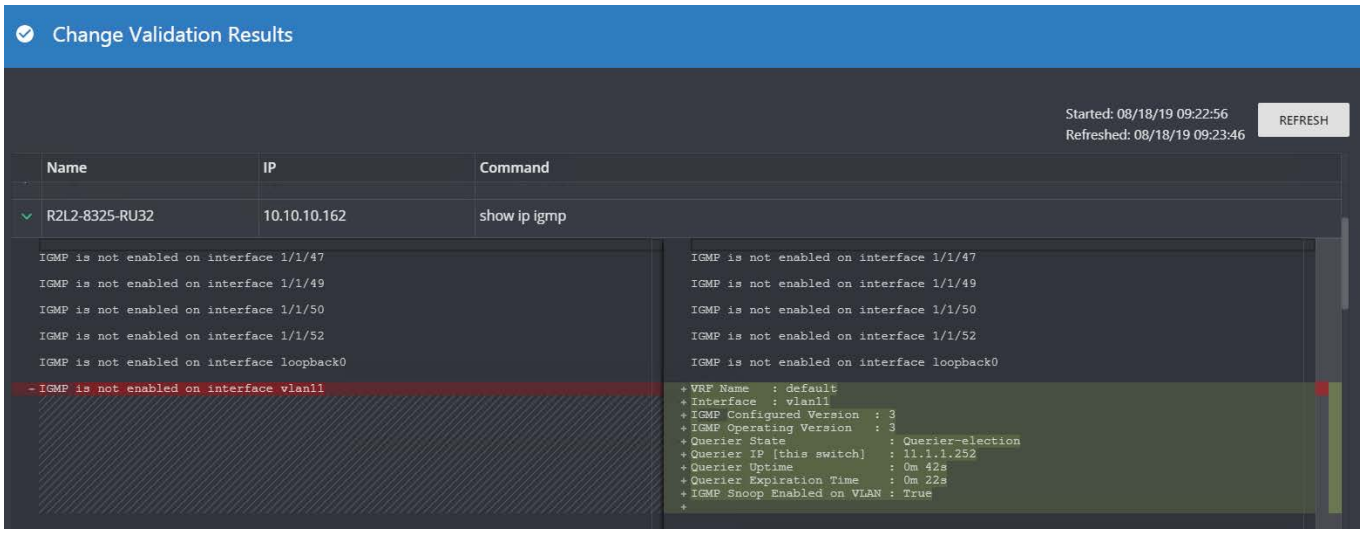

The image below is of 2 Hosts (Rack1 on the left and Rack 2 on the right). The host in Rack 2 is using a Multicast generator to source traffic. The host in Rack 1 is interested in listening to the MC traffic. Because MC is not yet configured end to end rack 1 is unable to receive the streams. This is expected at this stage.

Figure 59. Multicast source in Rack 2 not seen by Rack 1

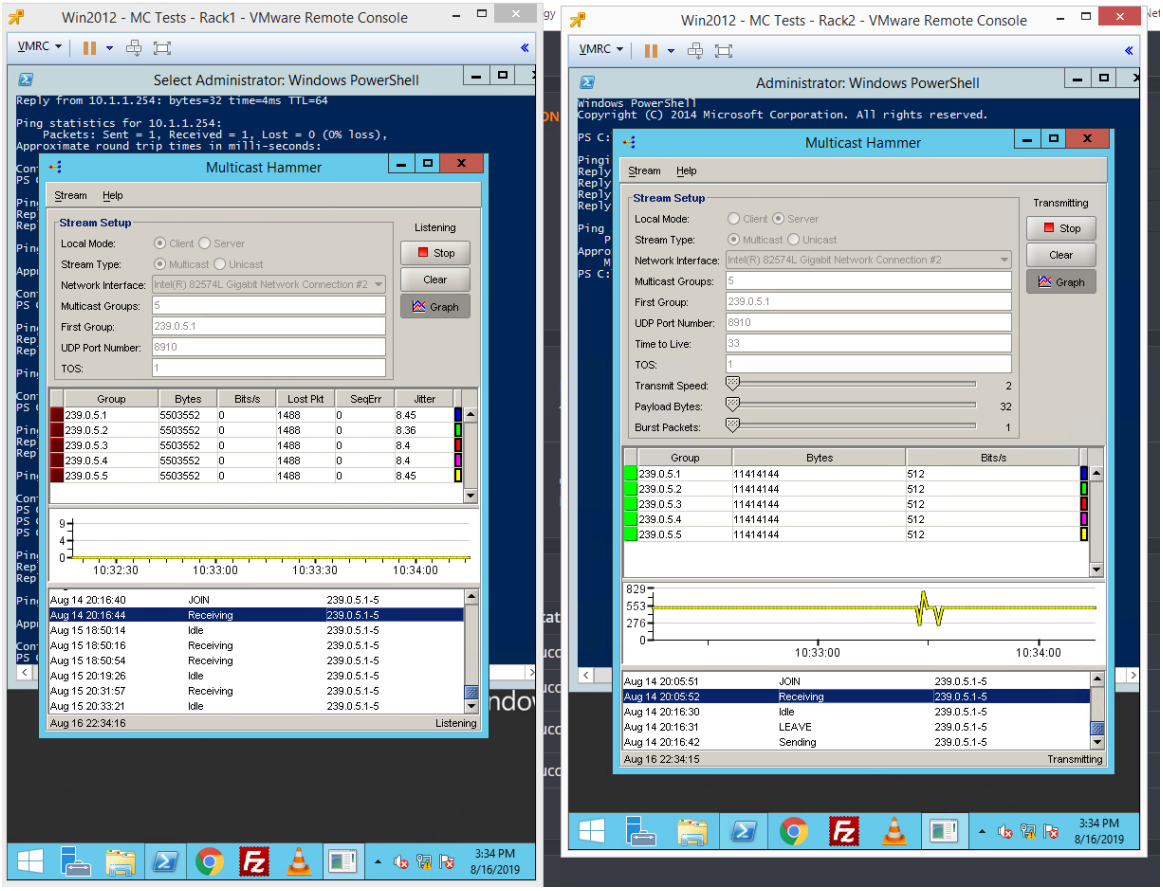

NetEdit Plan For Multicast Configuration on Spine Switches

Create a Multicast Leaf Configuration plan for Spine 1 & Spine 2 switches in NetEdit [Devices -> select both spines -> Action -> Edit Running Config] and name the plan.

Give the plan a name and “Create”.

Figure 60. Selecting Spines for MC configuration

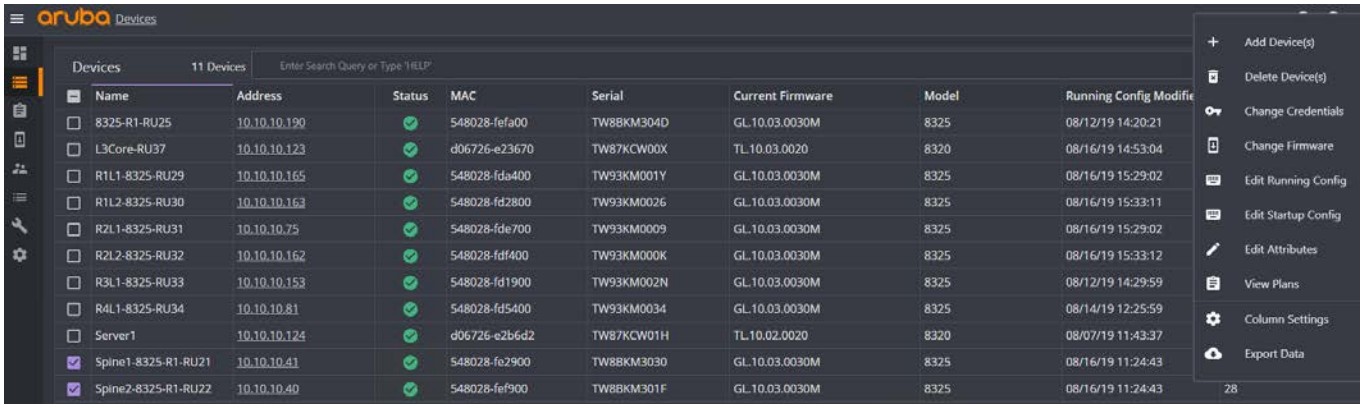

Add the following configuration.

```
interface 1/1/1
    ip pim-sparse enable
interface 1/1/2
    ip pim-sparse enable
interface 1/1/3
    ip pim-sparse enable
interface 1/1/4
    ip pim-sparse enable
interface loopback 0
    ip pim-sparse enable
router pim
    enable
    rp-address 1.1.1.1
```

Figure 61. PIM configuration added globally to both spines

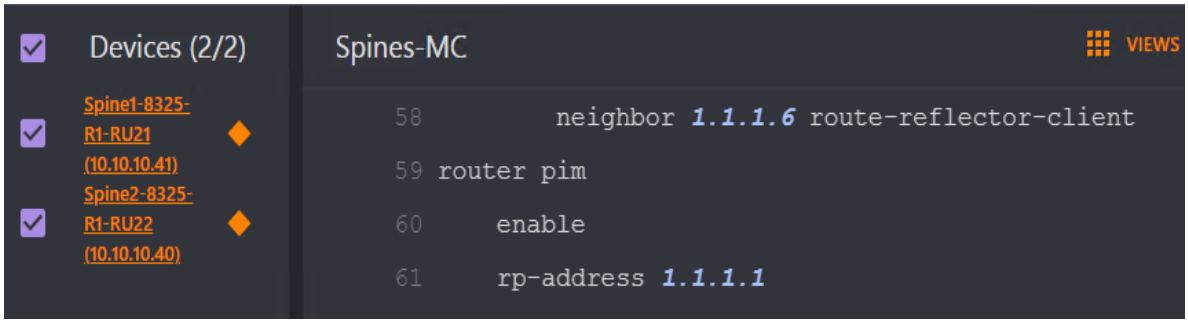

Figure 62. PIM Sparse configuration added to fabric facing interfaces

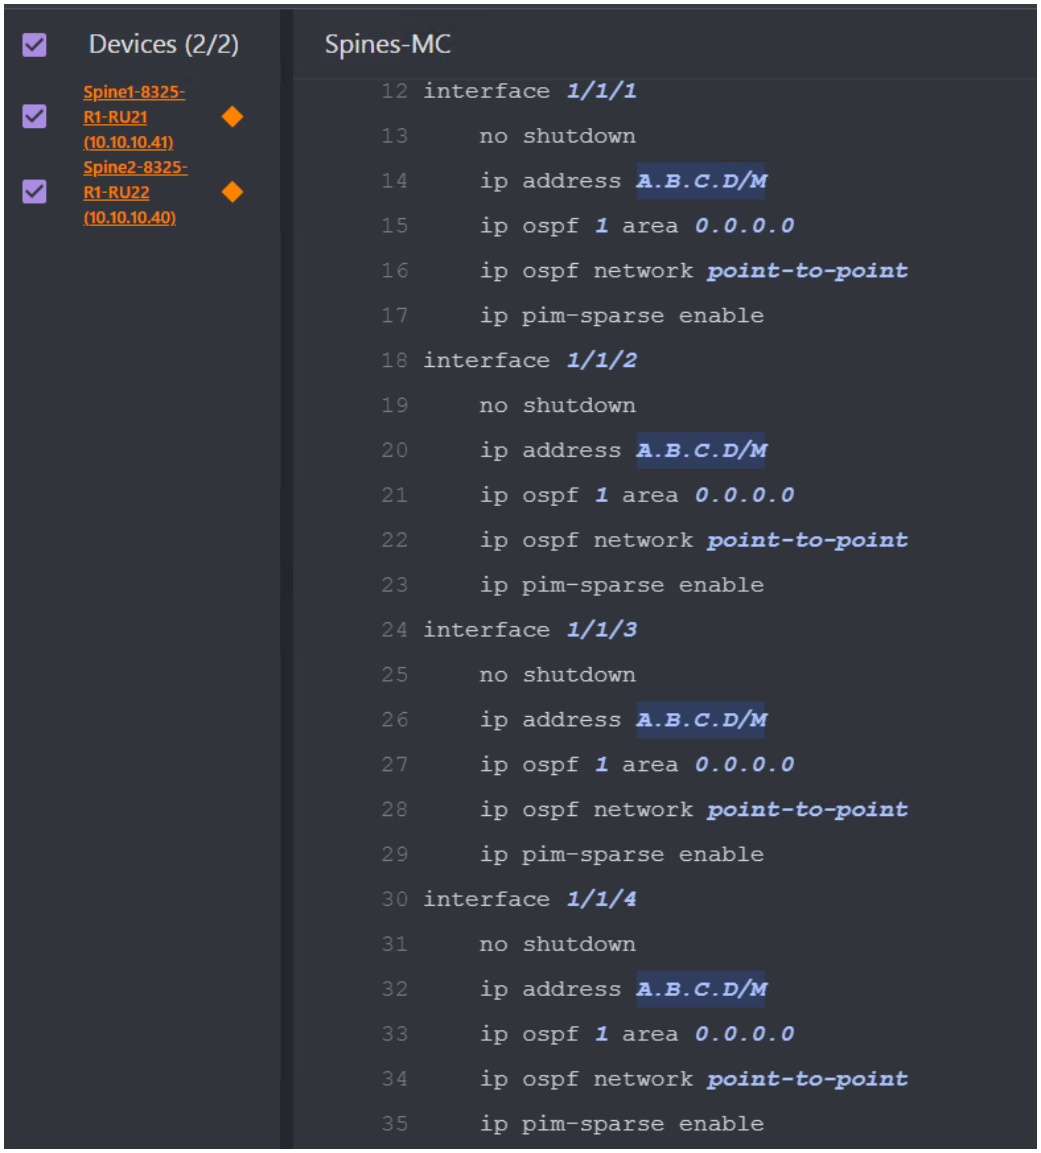

Now select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to PIM status.

Figure 63. Spine 2 showing PIM is now configured

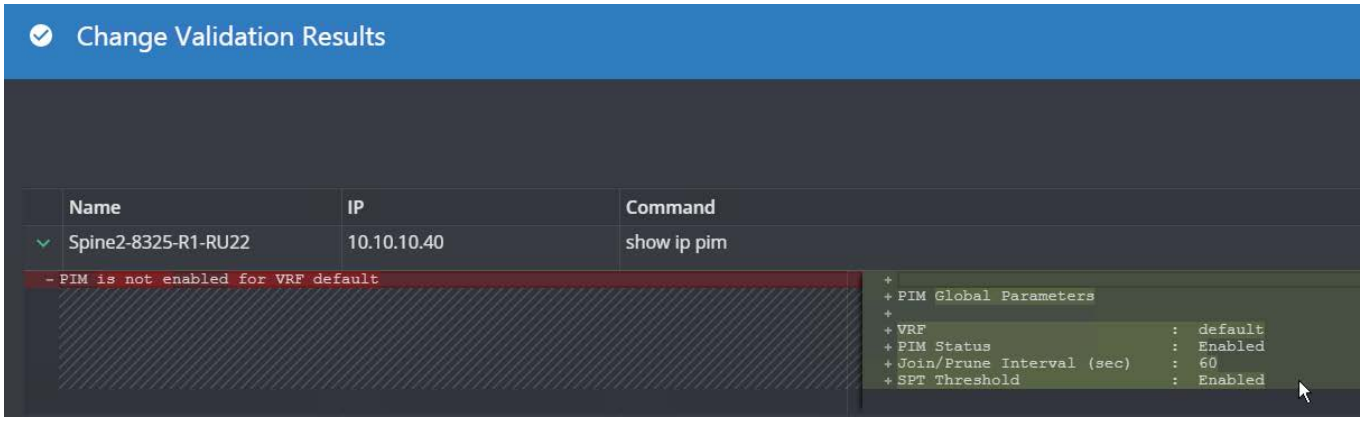

Figure 64. Spine 2 is now showing PIM neighbors

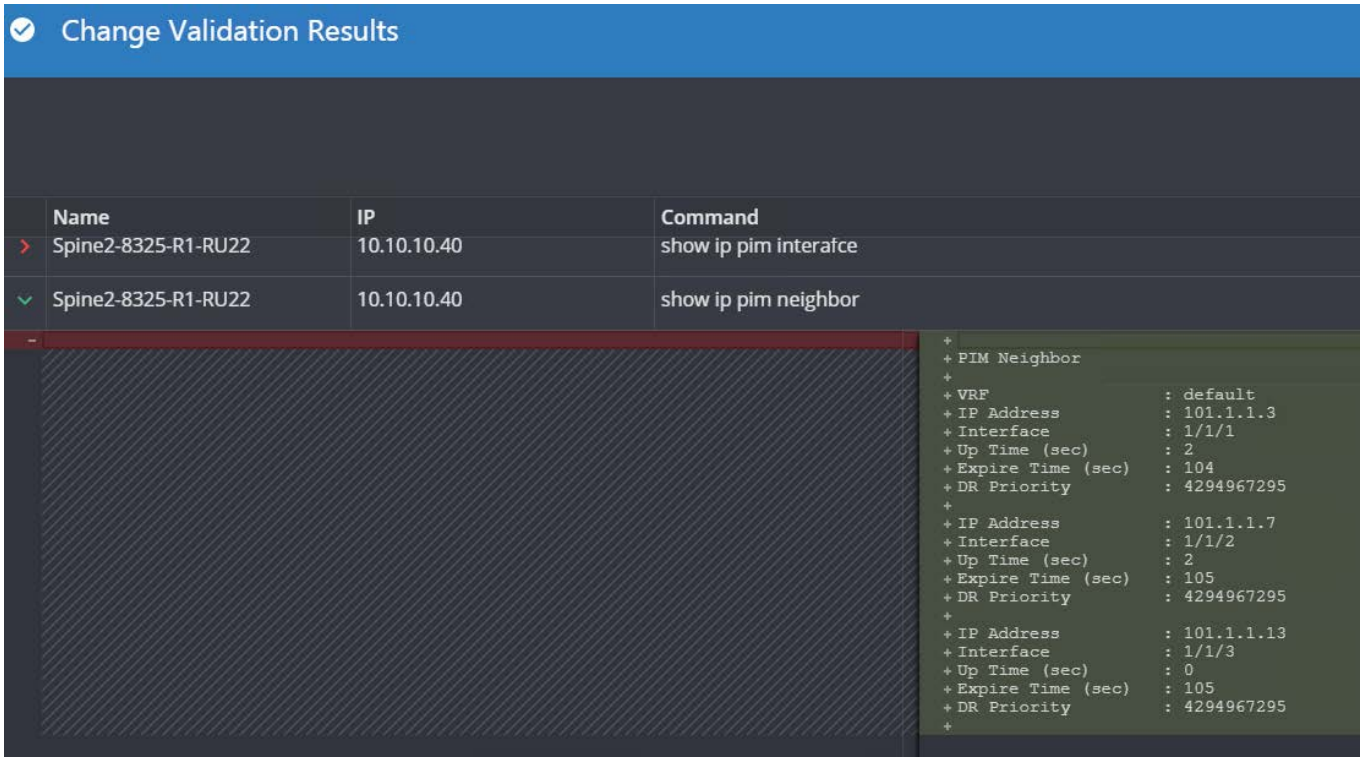

Figure 65. Spine 2 showing which interfaces PIM has been enabled on

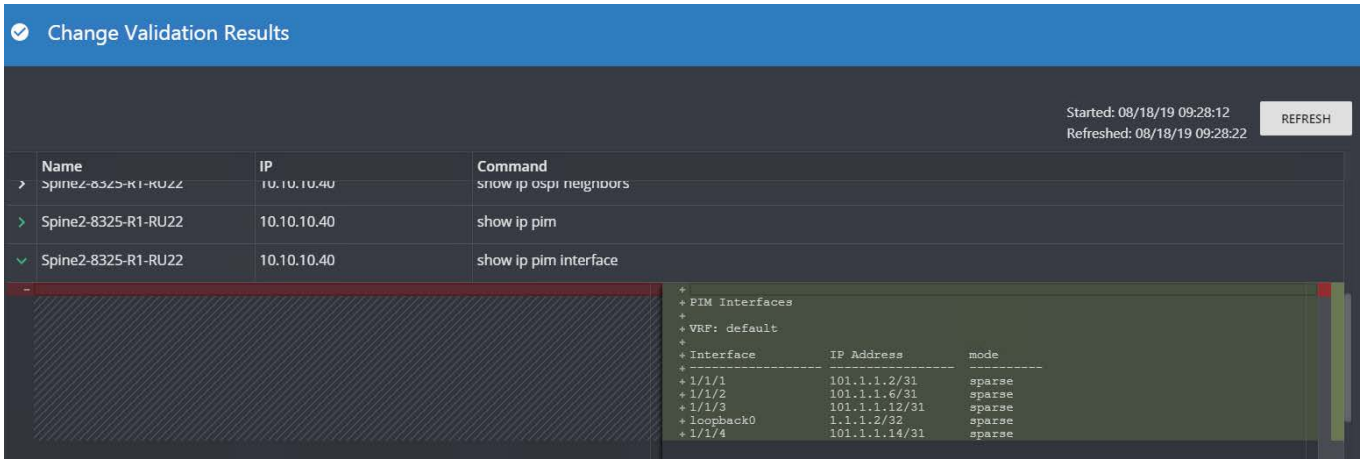

Figure 66. Spine 1 showing which interfaces PIM has been enabled on

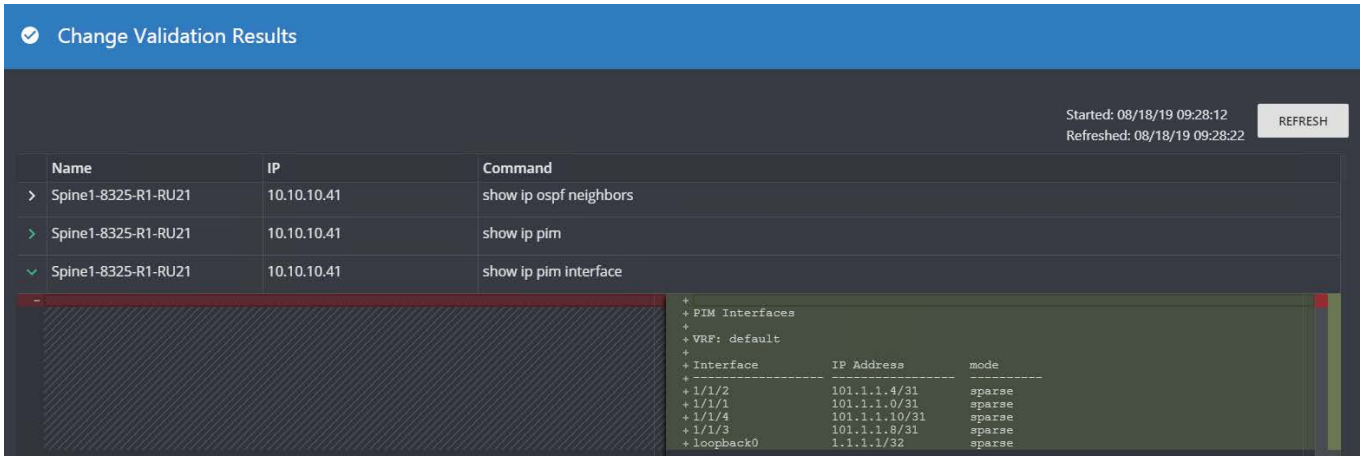

Figure 67. Spine 1 showing PIM is enabled with neighbors

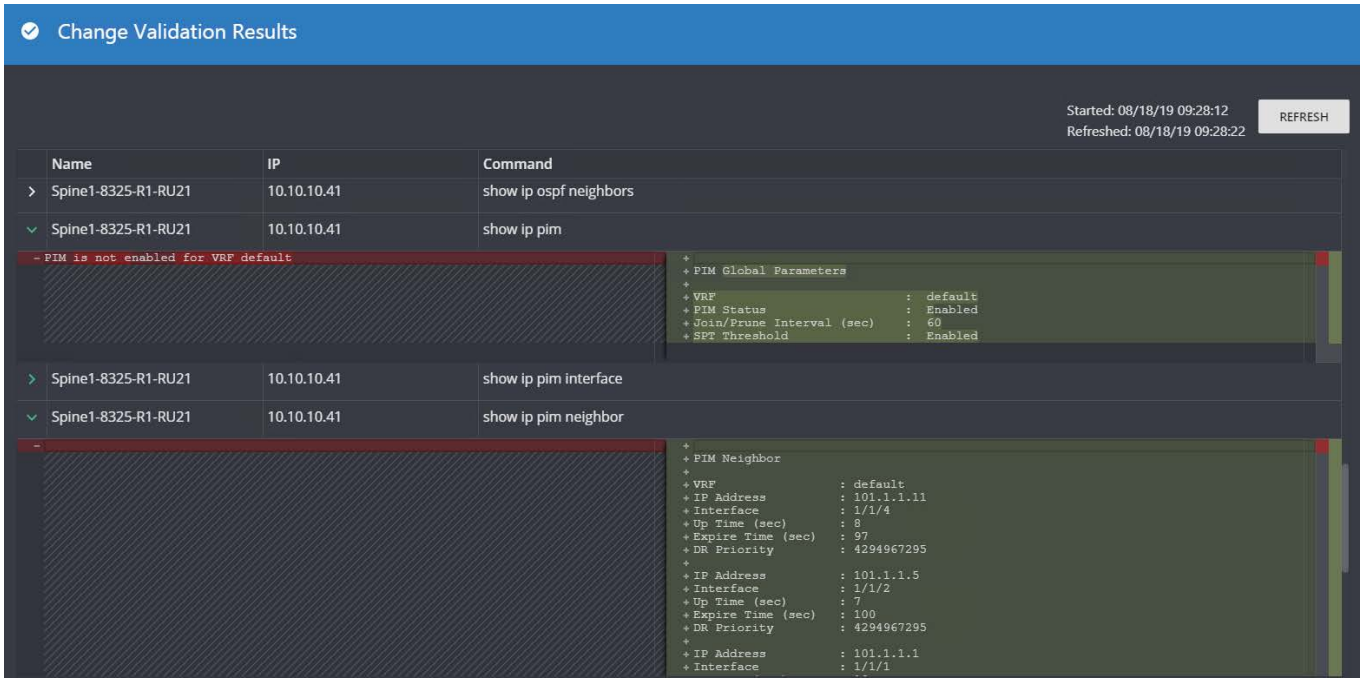

Figure 68. Spine 1 showing Multicast Routes for the 5 groups, and showing which interfaces it is being forwarded on

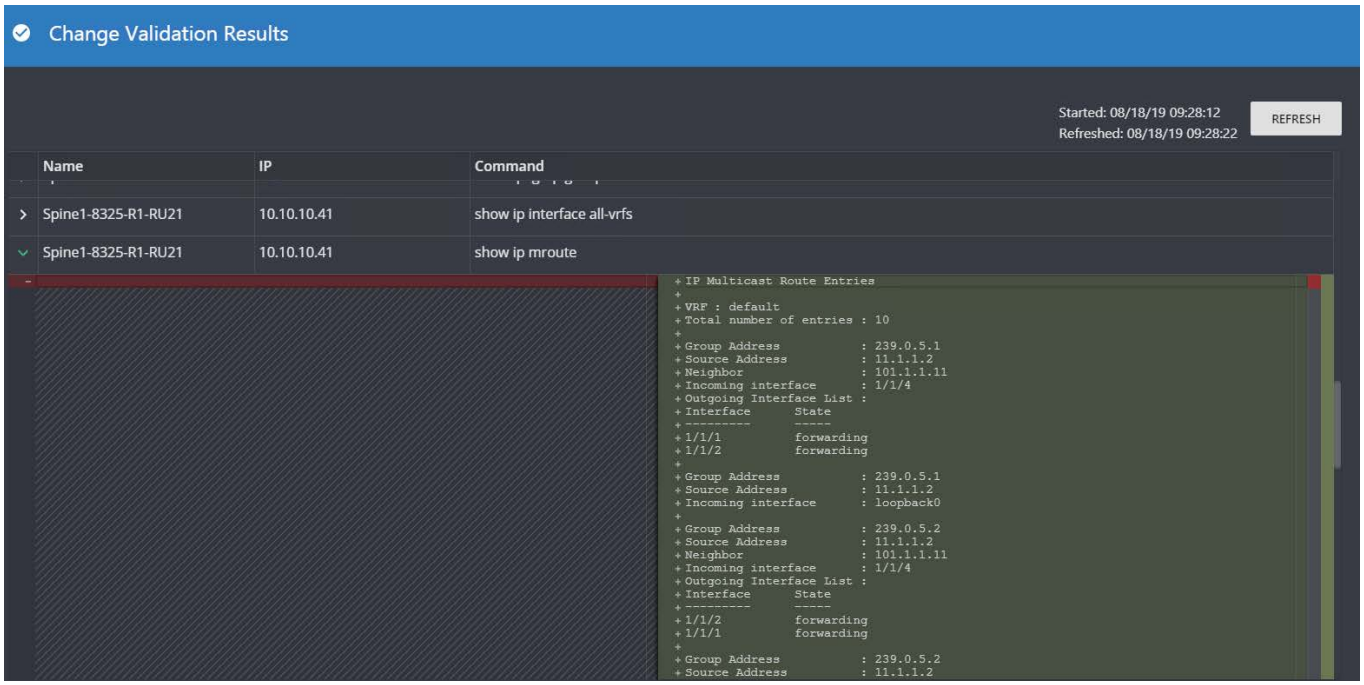

The image below is now showing that the host in Rack1 is receiving the MC streams.

Figure 69. Multicast source and receivers in Spine/Leaf now receiving multicast traffic

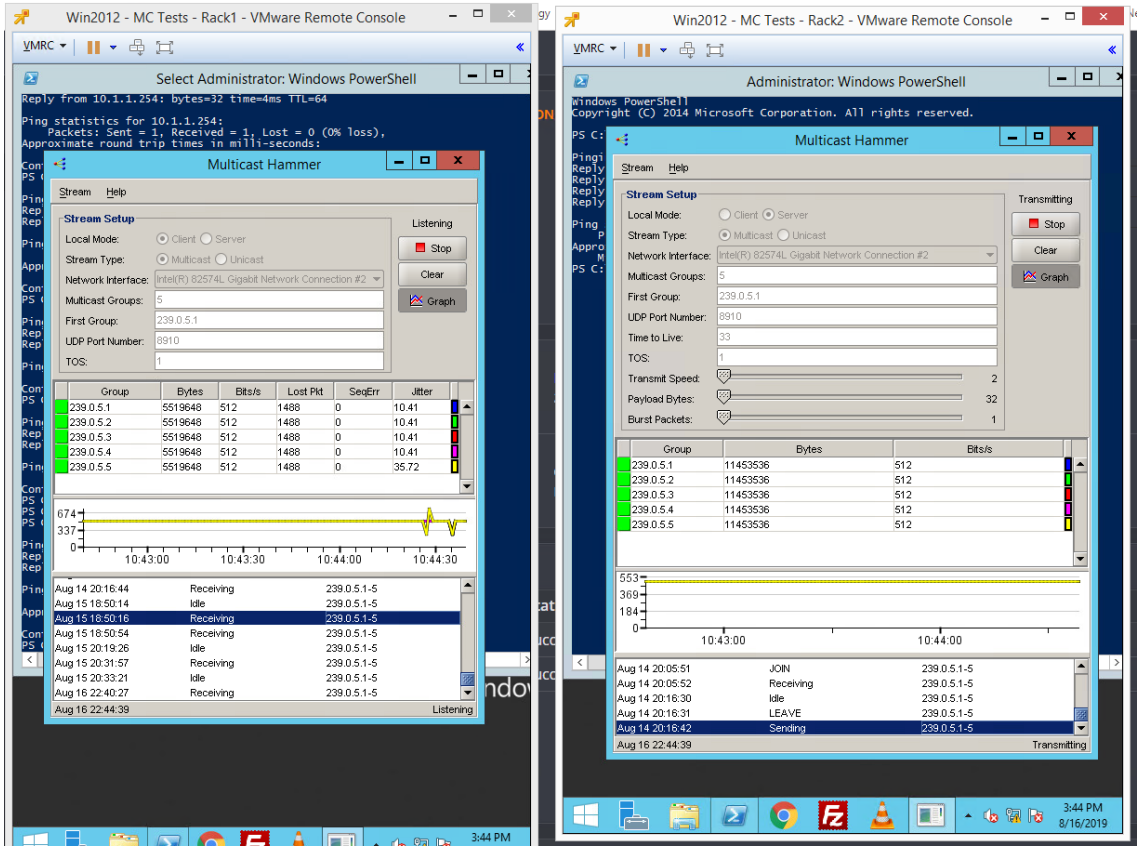

NetEdit Plan For Multicast Configuration to Enable Connections to L3Core

Create a Multicast Leaf Configuration plan for Spine 1 & L3Core switches in NetEdit [Devices -> select spine 1 & L3Core -> Action -> Edit Running Config] and name the plan.

Give the plan a name and “Create”.

Figure 70. Selecting Spine 1 and L3Core switches for inter PIM configuration

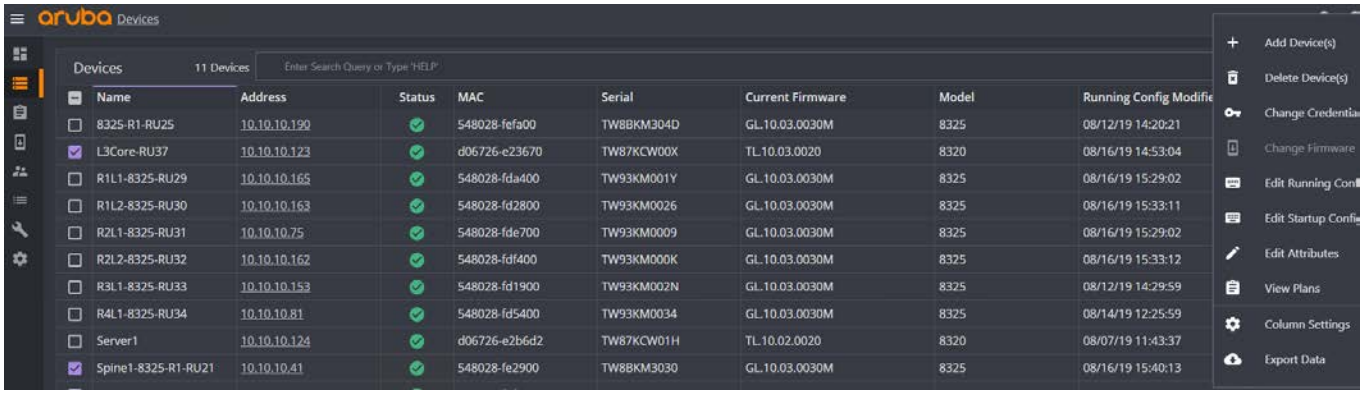

Choose the appropriate switches from the left panel and apply the required configurations.

 # L3Core

```
router pim
    enable
    rp-address 1.1.1.7
router msdp
    enable
    ip msdp peer 1.1.1.1
        connect-source loopback0
        enable
vlan 31
    ip igmp snooping enable
interface 1/1/49
    ip pim-sparse enable
interface 1/1/50
    ip pim-sparse enable
interface loopback 0
    ip pim-sparse enable
interface vlan31
    ip pim-sparse enable
    ip igmp enable
```

 # Spine1

```
router msdp
    enable
    ip msdp peer 1.1.1.7
        connect-source loopback0
        enable
```

Figure 71. L3Core showing new Global PIM and MSDN configurations

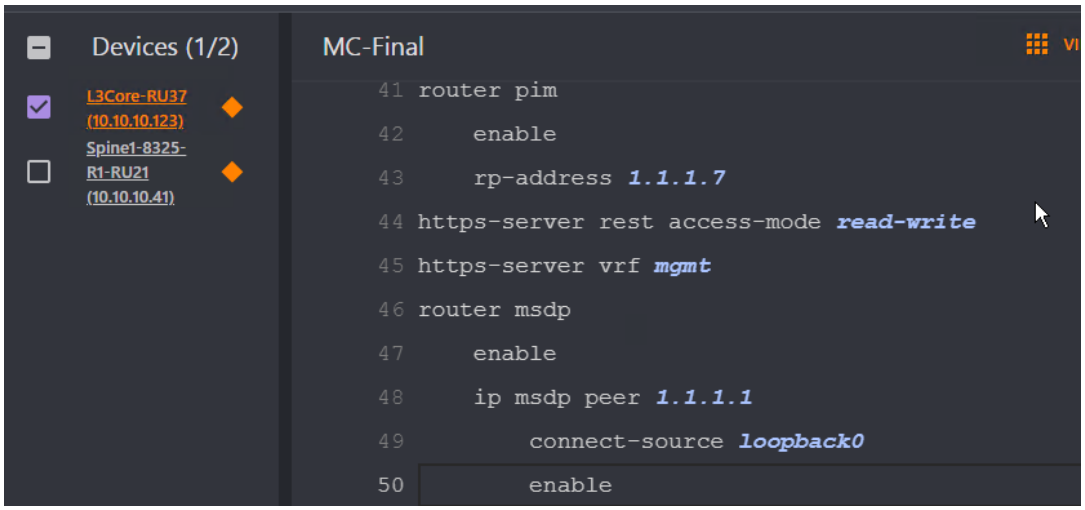

Figure 72. Spine 1 MSDN configuration

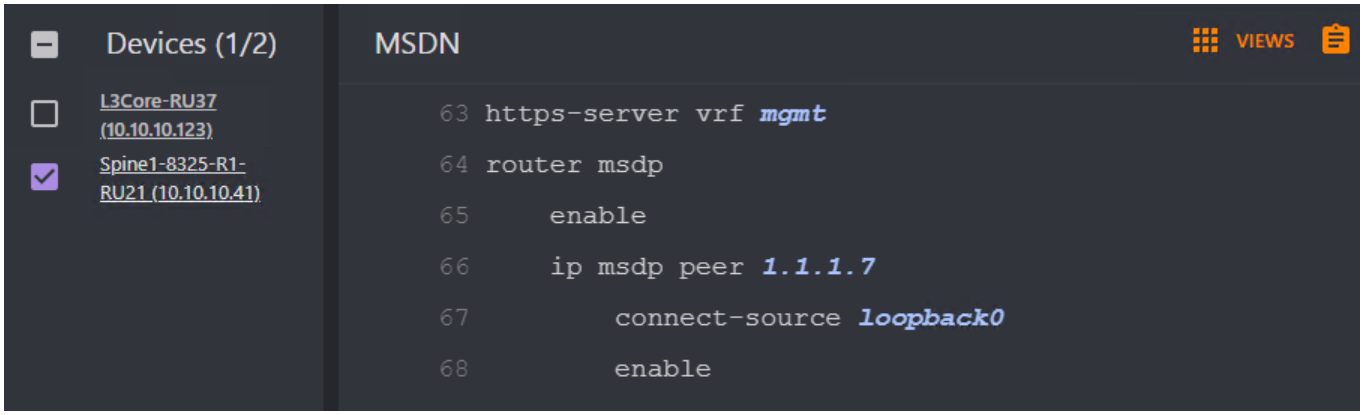

Figure 73. L3Core showing new Global PIM configuration on Interfaces

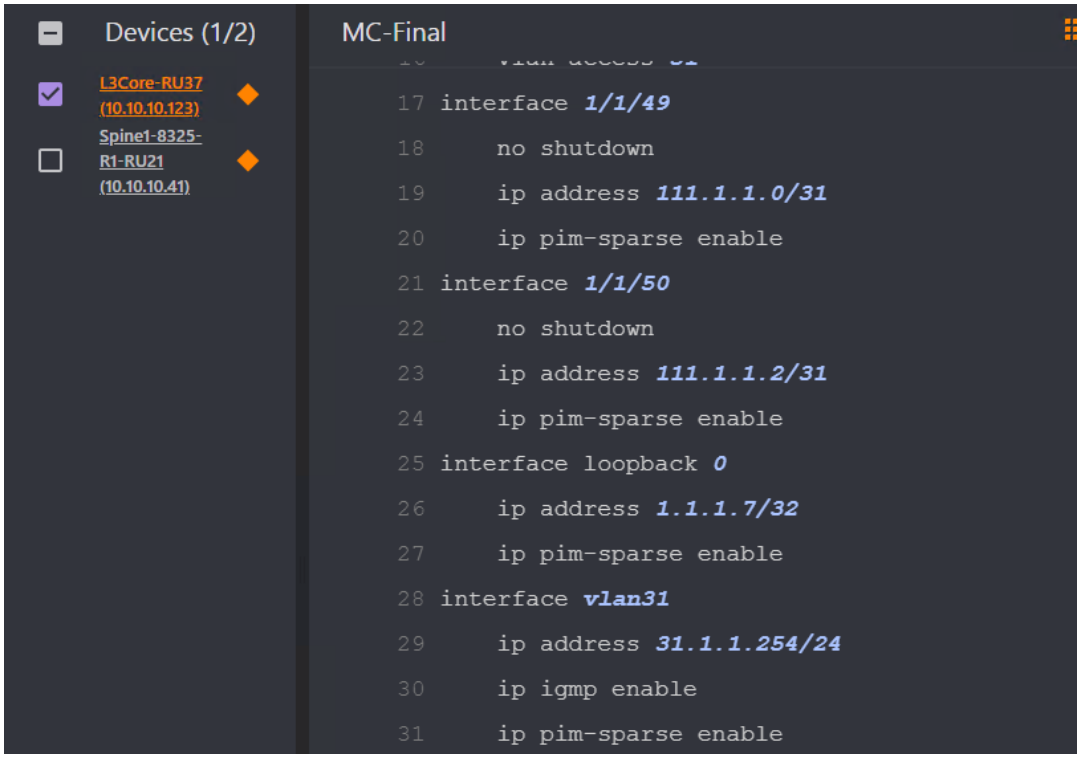

Figure 74. L3Core showing new IGMP snooping configuration

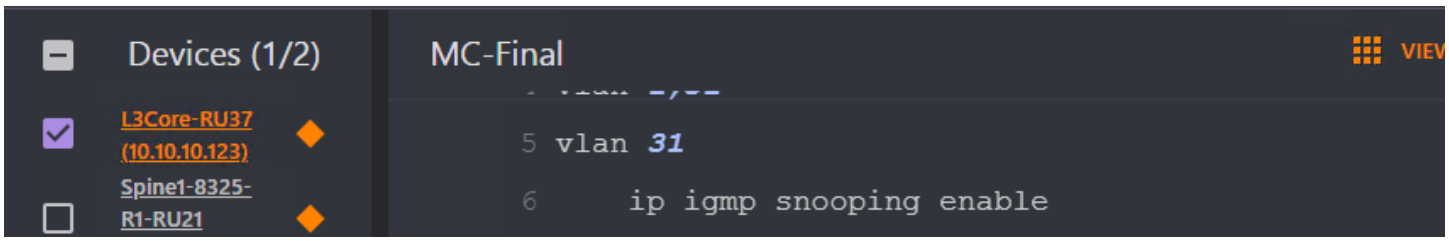

Finally, select “RETURN TO PLAN” -> “DEPLOY” to push down configs.

Verify Change Validation

You can click on “Change Validation” to verify that MSDP peering is up, that the L3Core can see Multicast Routes.

Figure 75. L3Core showing MSDN peering and that it is seeing Multicast routes from peer

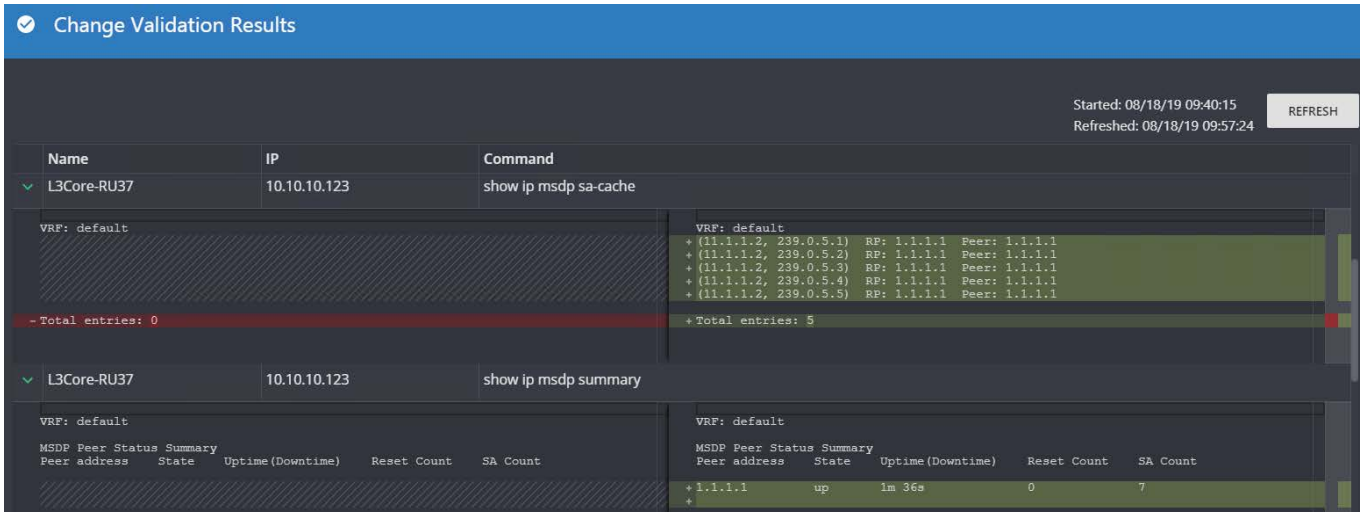

Figure 76. L3Core showing Multicast Routes from neighboring PIM domain

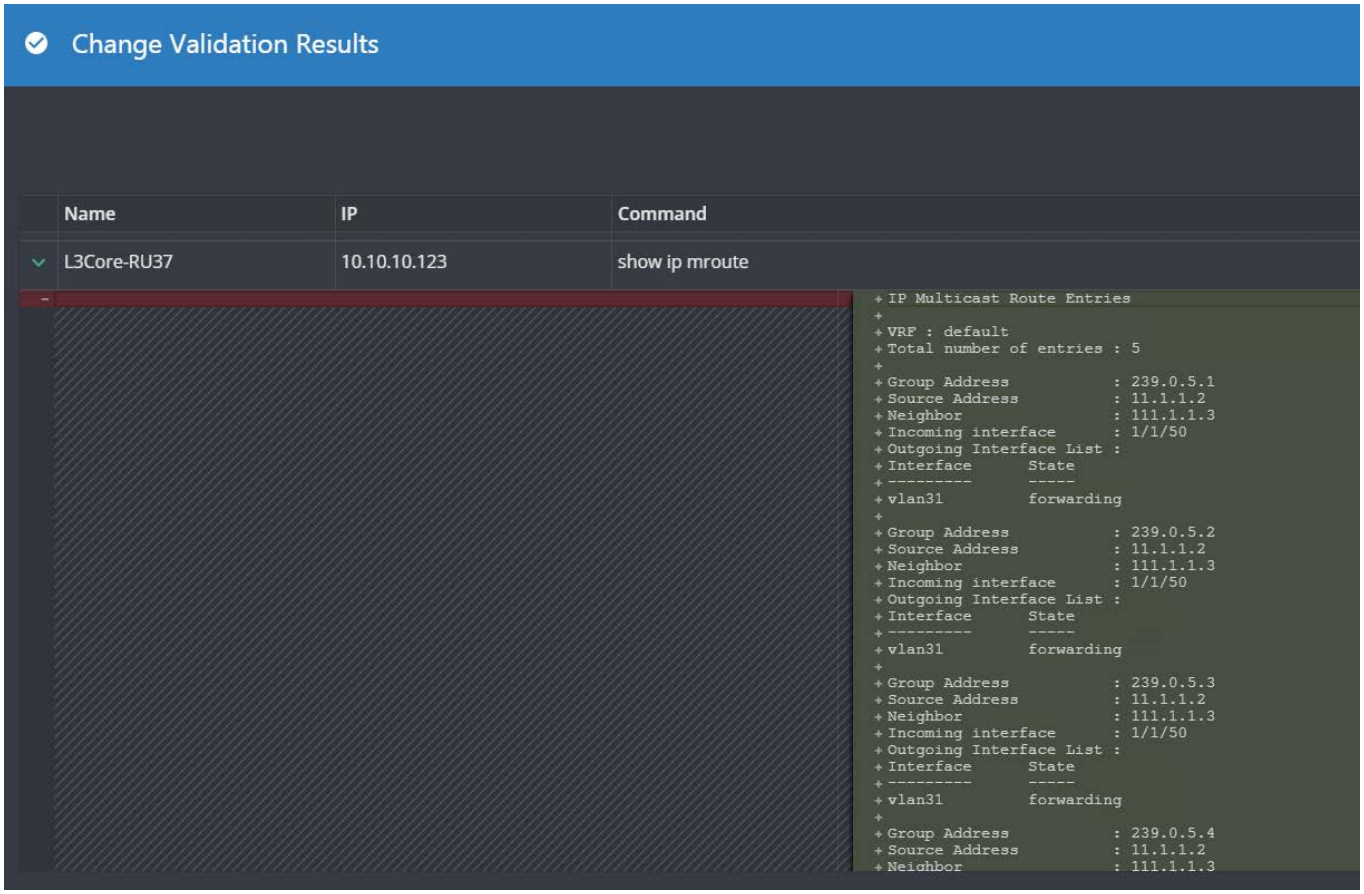

The image below is now showing that the host in L3Core is now receiving the MC streams.

Figure 77. Multicast source and receivers seeing multicast traffic between domains

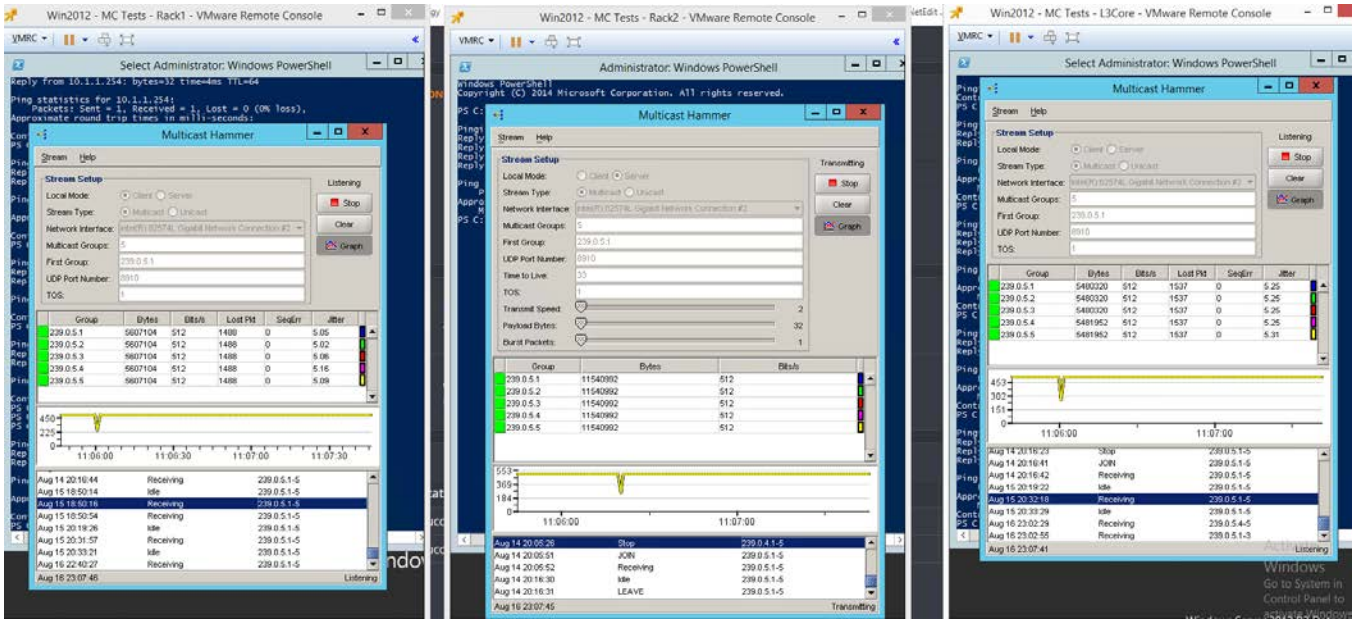

After validation, if things are looking good as the above screen shots show, you can choose to “COMMIT” to save the desired configs to the startup configuration or you can click “ROLLBACK” to revert configs before the configs were deployed and make further desired changes.