Juniper Networks | Engineering Simplicity

Network Configuration Example

Junos OS Upgrade for EVPN VXLAN Network

Published

2023-04-27

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Network Configuration Example Junos OS Upgrade for EVPN VXLAN Network

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

CHAPTER 1 EVPN VXLAN Upgrade Overview

About This Configuration Example

Overview

Use this network configuration example (NCE) to upgrade Junos OS for Ethernet VPN (EVPN) Virtual Extensible LAN protocol (VXLAN) network.

The upgrade procedures in this document are based on recommended configurations in Juniper Data Center EVPN-VXLAN Fabric Architecture Guide. Both ERB/Collapsed Spine and CRB deployment are covered. For more information, see Data Center EVPN VXLAN Fabric Architecture Guide.

The document covers the following platforms:

- QFX10002, QFX10008, QFX10016
- QFX5100
- QFX5110
- QFX5120-48Y
- No VC/VCF
- Multi-home ESI connection to host

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation.
Send your comments to design-center-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

CHAPTER 2 Plan Upgrade for EVPN VXLAN

Planning the Upgrade

This section covers the guidelines for planning the upgrade.

- Check the current bandwidth usage of all network links. We recommend you add links in advance to avoid bandwidth over-subscription when traffic is diverted during the device upgrade. In particular, the DCI connections and the links to the firewall or load balancers.
- This upgrade procedure expects the devices to be managed via out-of-band and via the management interface. When we shut down all ports, this action breaks management access, if currently done in-band. Check and add out-of-band management, if needed.
- There are some CLI show command that are executed, and information is collected. It is recommended to automate the CLI commands in scripts. Ensure to store all the information on the server and use tools that can quickly search and compare certain information from all collected information.
- Identify and design some traffic flows that can validate the network connectivity and run in background during the upgrade. Ping and traceroute are commonly used for this purpose.
- Plan sufficient time for the maintenance of the window. Multiple maintenance windows might be needed for large deployment. Always upgrade one device at a time at initial phase. After a few successful upgrades and getting familiar with the procedure, upgrade can be done in batches with multiple single members of the leaf pairs at a time for large deployment.
- Schedule the change with all teams that are affected by the changes.

Pre-Upgrade and Post-Upgrade Health Checking

This section covers the guidelines for planning the upgrade.

- We highly recommended performing health checks. Record this information so it can be compared between pre-upgrade to post-upgrade information in the event of an issue.
- Before the upgrade procedure listed in the document, perform customized checking and preparation procedures for the planned upgrade device after finishing the health checking.
- Check user traffic flow as expected without loss before and after the upgrade.
- Backup all devices and server' configurations before and after the upgrade.
- Collection of detail information and check the system is in a healthy state before and after the upgrade as follows:
    - Check syslog for any failures and errors.
      ```
      show log messages | no-more
      show log chassisd | no-more
      ```
    - Check alarms and core-dump on the system.
      ```
      show chassis alarms | no-more
      show system alarms | no-more
      show system core-dumps | no-more
      ```

- Check RE/FPC/PIC status and interfaces status as expected.

```
show chassis fabric summary
show chassis fabric fpcs
show chassis hardware detail | no-more
show chassis fpc detail | no-more
show chassis fpc pic-status | no-more
show chassis environment | no-more
show chassis routing-engine | no-more
show interfaces descriptions | match up | no-more
show interfaces descriptions | match down | no-more
show interface xe-* | match "physical|rate"
show interface et-* | match "physical|rate"
```

- Check routing table and forwarding table as expected.

```
show system processes extensive | no-more
show krt queue | no-more
show route summary | no-more
show route forwarding-table summary | no-more
show arp no-resolve expiration-time | no-more
```

- Check and take detail information of EVPN related information.

```
show arp no-resolve | count
show ethernet-switching table | count
show evpn instance
show evpn database
show ethernet-switching table
show ethernet-switching vxlan-tunnel-end-point remote summary
show ethernet-switching vxlan-tunnel-end-point esi
```

- Check system free storage to ensure enough storage for the new Junos OS image.
  Run `df -k /var/tmp` at shell mode for checking free space.
- Use `request system storage cleanup` command to free storage space on the devices.
  Always run `request system storage cleanup dry-run` first to check the proposing list of files for deletion.
- Clear the alarms and core-dumps before the upgrade.

```
clear system error
```

CHAPTER 3 EVPN VXLAN System Upgrade in ERB Leaf Node

System Upgrade in ERB Leaf Node

ERB Reference Network Deployment Overview

In this deployment:

- Lean spines, as shown in blue box in the figure, provide:
    - L3 forwarding for underlay routing
    - eBGP overlay
- ERB leaf, as shown in grey box in the figure, are the devices to be upgraded.
- End-device facing L2 access interfaces running LACP are shown with green lines.
- Network core facing L3 interfaces are shown with blue lines.

Figure 1: ERB Reference Network Deployment Topology

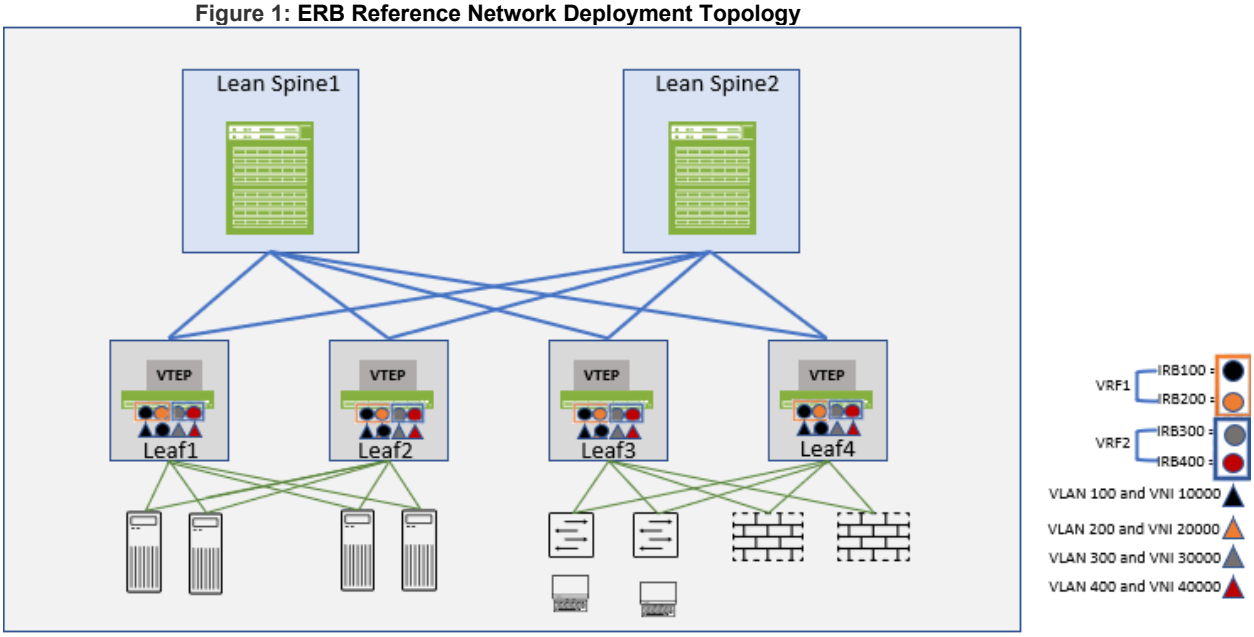

NOTE: Always perform pre-upgrade health check before next step. Upgrade one device at a time.

Step-by-Step Procedure

Follow these steps to upgrade in ERB reference network deployment.

1. Prepare for the upgrade.
   a. Backup device and server configurations:
   b. Clean up system storage.
      - `request system storage cleanup`
      - Always run `request system storage cleanup dry-run` first to check the proposing list of files for deletion.
   c. Copy Junos OS image to device `/var/tmp` directory.
   d. Enable syslog ("any" "info") to capture detailed logs:
      ```
      set system syslog file messages any info
      ```
      You can use this information for debugging any issues.
   e. Disable HA knobs configurations on device under upgrade (QFX10008 and QFX10016 only):
      ```
      deactivate chassis redundancy graceful-switchover
      ```
2. Disable all end-device facing access interfaces to divert traffic towards multi-homed peer devices.
   a. Verify all end-device facing interfaces are down.
      ```
      show interfaces descriptions | match up | no-more
      show interfaces descriptions | match down | no-more
      ```
   b. Verify no incoming traffic from network core (except control packets).
      ```
      monitor interface traffic
      ```
3. Disable all network core facing interfaces.
   a. Verify all interfaces are down.
      ```
      show interfaces descriptions | match up | no-more
      show interfaces descriptions | match down | no-more
      ```
   b. Ensure that the following configuration does not exist for each vlan. If the configuration is existing, remove the configuration.
      ```
      set vlans <vlan-name> vxlan ingress-node-replication
      ```
      NOTE: Do not remove the config under protocol evpn.
      ```
      set protocols evpn multicast-mode ingress-node-replication
      ```
4. Upgrade device with new image and reboot.
   a. Upgrade backup RE first then upgrade active RE with reboot:
      1. Run `show chassis hardware` to check if system has the backup RE .
      2. Run `show chassis routing-engine` to determine the primary and backup RE.

         In normal situation, RE1 is the backup RE while RE0 is the primary.
      3. Run `request system software add /var/tmp/new_junos_image.tgz re1 force-host reboot` to upgrade backup RE first and wait until backup RE reboots.

4. Run `show chassis routing-engine` to check RE status.
5. Run `request system software add /var/tmp/new_junos_image.tgz force-host reboot` to proceed with primary RE upgrade.

b. Wait and verify the device is up.

c. Verify no process cores.
```
show system core-dumps
```

d. Verify no additional system and chassis alarms.
```
show chassis alarms | no-more
show system alarms | no-more
```

5. Enable all core-facing interfaces.

a. Wait and verify all underlay routing are up.

b. Wait and verify BGP neighbors are established. Wait for BGP route update to be finished.
Run `show bgp summary` and check OutQ close to 0 for all neighbors.

c. Wait and verify IRB interfaces are up.
```
show interfaces irb
```

d. Verify VXLAN tunnel to remote VTEPs are up.
```
show ethernet-switching vxlan-tunnel-end-point remote
```

e. Verify remote mac learning correctly in rpd and l2ald.
```
show evpn instance
show ethernet-switching global-mac-count
show ethernet-switching table summary
```

6. Enable all end-device facing access interfaces.

a. Hosts might have small amount of traffic loss during this period.

b. Wait and verify all multi-home ESIs are up and established with remote peer.
```
show ethernet-switching vxlan-tunnel-end-point esi
```

c. Wait and verify all user traffic is resumed as normal.

> **NOTE:** Always perform post-upgrade health check before next step.

d. Contact customer support for alarms and core-dumps during and after the upgrade.

- Provide the syslog file messages and core files.
- Provide RSI info capture on the devices.
  This is applicable from Junos OS 21.2R3-S3, Junos OS 21.4R3, and later releases. The execution time for this command varies based on the scale of the devices.

```
request support information evpn-vxlan | tee /var/tmp/rsi.txt
```

e. Contact customer support for rolling back the Junos OS software, if upgrade fails
```
request system software rollback
request system reboot
```

7. Post-upgrade cleanup actions include:

    a. Clean up system storage.

    b. Remove the Junos OS image.

    c. Enable back HA knobs configurations on device after upgrade (QFX10008, QFX10016 only).

    ```
    activate chassis redundancy graceful-switchover
    ```

    d. Restore syslog settings back to original.

CHAPTER 4 EVPN VXLAN System Upgrade in CRB Deployment

System Upgrade in CRB Deployment

This procedure lists two steps in sequences.

> **NOTE:** We recommend that you upgrade all the leaf devices first, and then proceed with the upgrade on Spine devices.

CRB Reference Network Deployment

In the deployment, all links in the topology:

- End-device facing L2 access interfaces running LACP are shown with green lines.
- Network core facing L3 interfaces are shown with blue lines.
- Spines, as shown in blue boxes, provide EVPN VXLAN L3 gateway functionality with IRB routing. Spines also perform iBGP router reflector functionality.
- Leaf, as shown in grey boxes, provides L2 gateway functionality for MH ESI connectivity.

Figure 2: CRB Reference Network Deployment Topology

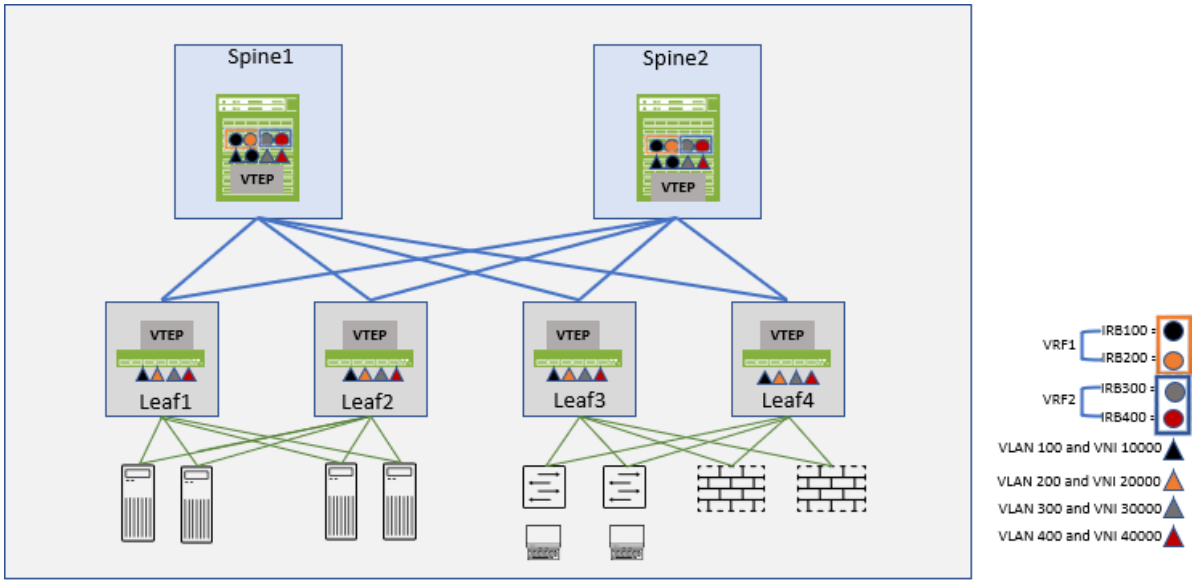

CHAPTER 5 EVPN VXLAN Upgrade for CRB Deployment - Leaf Devices

CRB Deployment - Leaf Devices Upgrade

NOTE: Always perform pre-upgrade health check before the next step.

Step-by-Step Procedure

Follow these steps to upgrade in CRB reference network deployment.

1. Preparation for the upgrade.
   a. Backup device and server configurations.
   b. Clean up system storage.
      - `request system storage cleanup`
      - Always run `request system storage cleanup dry-run` first to check the proposing list of files for deletion.
   c. Copy Junos OS image to device `/var/tmp` directory.
   d. Enable syslog ("any" "info") to capture detailed logs:
      ```
      set system syslog file messages any info
      ```
      You can use this information for debugging any issues.
   e. Disable HA knobs configurations on device under upgrade (QFX10008 and QFX10016 only).
      ```
      deactivate chassis redundancy graceful-switchover
      ```
2. Disable all end-device facing access interfaces to divert traffic for local connected devices.
   a. Verify all end-device facing interfaces are down.
      ```
      show interfaces descriptions | match up | no-more
      show interfaces descriptions | match down | no-more
      ```
   b. Verify that there is no incoming traffic from network core (except control packet).
      ```
      monitor interface traffic
      ```
   c. Ensure that the following configuration does not exist for each vlan. If the configuration is existing, remove the configuration.
      ```
      set vlans <vlan-name> vxlan ingress-node-replication
      ```
      NOTE: Do not remove the config under protocol evpn.
      ```
      set protocols evpn multicast-mode ingress-node-replication
      ```
3. Upgrade device with new image and reboot.
   a. Upgrade backup RE first then upgrade active RE with reboot.
      1. Run `show chassis hardware` to check if system has the backup RE .

2. Run `show chassis routing-engine` to determine the primary and backup RE.

   In normal situation, RE1 is the backup RE while RE0 is the primary.

3. Run `request system software add /var/tmp/new_junos_image.tgz re1 force-host reboot` to upgrade backup RE first and wait until backup RE reboots.
4. Run `show chassis routing-engine` to check RE status.
5. Run `request system software add /var/tmp/new_junos_image.tgz force-host reboot` to proceed with primary RE upgrade.

b. Wait and verify the device is up.

c. Verify that there are no process cores.

```
show system core-dumps
```

d. Verify that there are no additional system and chassis alarms.

```
show chassis alarms | no-more
show system alarms | no-more
```

4. Enable the network core facing interfaces.

   a. Hosts might have small amount of duplication packets received and/or traffic lost during this period.

   b. Wait and verify all underlay routing is up.

   c. Wait and verify all BGP neighbors are established. Wait for BGP route update to be finished.

      Run `show bgp summary` and check OutQ close to 0 for all neighbors.

   d. Verify VXLAN tunnel to remote VTEPs are up.

```
show ethernet-switching vxlan-tunnel-end-point remote
```

   e. Verify remote mac learning correctly in rpd and l2ald.

```
show evpn instance
show ethernet-switching global-mac-count
show ethernet-switching table summary
```

5. Enable all end-device facing access interfaces.

   a. Hosts might have small amount of duplication packets received and/or traffic lost during this period.

   b. Wait and verify all interfaces are up.

```
show interfaces descriptions | match up | no-more
show interfaces descriptions | match down | no-more
```

   c. Wait and verify all multi-home ESIs are up and established with remote peer.

```
show ethernet-switching vxlan-tunnel-end-point esi
```

   d. Wait and verify all user traffic is resumed as normal.

> **NOTE:** Always perform post-upgrade health check before next step.

e. Contact customer support for alarms and core-dumps during and after the upgrade.

- Provide the syslog file “messages” and core files.
- Provide RSI info capture on the devices. This is applicable from Junos OS 21.2R3-S3, Junos OS 21.4R3, and later releases. The execution time for this command varies based on the scale of the devices.

```
request support information evpn-vxlan | tee /var/tmp/rsi.txt
```

f. Contact customer support for rolling back the Junos OS software, if upgrade fails

```
request system software rollback
request system reboot
```

6. Post-upgrade cleanup actions include:

a. Clean up system storage.

b. Remove the Junos OS image.

c. Enable back HA knobs configurations on device after upgrade (QFX10008 and QFX10016 only).

```
activate chassis redundancy graceful-switchover
```

d. Restore syslog settings back to originals.

CHAPTER 6 EVPN VXLAN Upgrade for CRB Deployment - Spine Devices

CRB Deployment - Spine Devices Upgrade

NOTE: Always perform pre-upgrade health check before the next step.

Step-by-Step Procedure

Follow these steps to upgrade CRB deployment for spine devices.

1. Prepare for the upgrade.
   a. Backup device and server configurations.
   b. Clean up system storage.
      - `request system storage cleanup`
      - Always run `request system storage cleanup dry-run` first to check the proposing list of files for deletion.
   c. Copy Junos OS image to device `/var/tmp` directory.
   d. Enable syslog ("any" "info") to capture detailed logs:
      ```
      set system syslog file messages any info
      ```
      You can use this information for debugging any issues.
   e. Disable HA knobs configurations on device under upgrade (QFX10008 and QFX10016 only).
      ```
      deactivate chassis redundancy graceful-switchover
      ```
2. Disable all access interface (if there is any) to divert traffic for local connected devices.
   a. Verify all end-device facing interfaces are down.
      ```
      show interfaces descriptions | match up | no-more
      show interfaces descriptions | match down | no-more
      ```
3. Disable all WAN/DCI/L3 interfaces and interfaces connected to all leafs.
   a. Verify that there is no incoming traffic from network core.
      ```
      monitor interface traffic
      ```
   b. Verify all user traffic divert to other spine and user traffic converged.
      Small amount of traffic loss is expected during this step.
   c. Ensure that the following configuration does not exist for each vlan. If the configuration is existing, remove the configuration.
      ```
      set vlans <vlan-name> vxlan ingress-node-replication
      ```

NOTE: Do not remove the config under protocol evpn.
`set protocols evpn multicast-mode ingress-node-replication`

4. Ensure that the following configuration exists for each irb IFL. If the configuration is not existing, add the configuration.
   `set interfaces irb unit xxxx proxy-macip-advertisement`
5. Upgrade device with new image and reboot.
   a. Upgrade backup RE first then upgrade active RE with reboot.
      1. Run `show chassis hardware` to check if system has the backup RE .
      2. Run `show chassis routing-engine` to determine the primary and backup RE.

         In normal situation, RE1 is the backup RE while RE0 is the primary.
      3. Run `request system software add /var/tmp/new_junos_image.tgz re1 force-host reboot` to upgrade backup RE first and wait until backup RE reboots.
      4. Run `show chassis routing-engine` to check RE status.
      5. Run `request system software add /var/tmp/new_junos_image.tgz force-host reboot` to proceed with primary RE upgrade.
   b. Wait and verify device is up.
   c. Verify that there are no process cores.
      `show system core-dumps`
   d. Verify that there are no additional system and chassis alarms.
      `show chassis alarms | no-more`
      `show system alarms | no-more`
6. Enable all WAN/DCI/L3 interfaces and interfaces connected to all leafs.
   a. Wait and verify all underlay routing is up.
   b. Wait and verify BGP neighbors established. Wait for BGP route update to finish.
      Run `show bgp summary and check OutQ close to 0 for all neighbors.`
   c. Wait and verify IRB interfaces are up.
   d. Verify VXLAN tunnel to remote VTEPs are up.
      `show ethernet-switching vxlan-tunnel-end-point remote`
   e. Verify remote mac learning correctly in rpd and l2ald.
      `show evpn instance`
      `show ethernet-switching global-mac-count`
      `show ethernet-switching table summary`
   f. Hosts might have small amount of duplication packets received and traffic loss during this period.
7. Enable all end-device facing access interfaces (if there is any).
   a. Hosts might have small amount of duplication packets or traffic loss during this period.

b. Wait and verify all interfaces are up.

```
show interfaces descriptions | match up | no-more
show interfaces descriptions | match down | no-more
```

c. Wait and verify all multi-home ESIs are up and established with remote peer.

```
show ethernet-switching vxlan-tunnel-end-point esi
```

d. Wait and verify all user traffic is resumed as normal.

e. Contact customer support for alarms and core-dumps during and after the upgrade.

- Provide the syslog file “messages” and core files.
- Provide RSI info capture on the devices. This is applicable from Junos OS 21.2R3-S3, Junos OS 21.4R3, and later releases. The execution time for this command varies based on the scale of the devices.

```
request support information evpn-vxlan | tee /var/tmp/rsi.txt
```

f. Contact customer support for rolling back the Junos OS software, if upgrade fails

```
request system software rollback
request system reboot
```

8. Post-upgrade cleanup actions include:

a. Clean up the system storage.

b. Remove the Junos OS image.

c. Enable back HA knobs configurations on device after upgrade (QFX10008 and QFX10016 only).

```
activate chassis redundancy graceful-switchover
```

d. Restore syslog settings back to original.

CHAPTER 7 EVPN VXLAN Upgrade Considerations

Configure IGMP Snooping Enabled Devices

This section is only applicable to IGMP Snooping enabled devices running Junos OS releases earlier than 21.3 and plan to upgrade to Junos OS 21.3R1 or later releases.

For EVPN devices enabled with IGMP snooping or MLD snooping, an additional configuration needs to be added after the device boot up with new Junos OS image and before enable any interfaces back, to inter-op with other devices still running with Junos OS releases earlier than 21.3R1.

See Junos OS EVPN User Guide for more information.

To verify whether IGMP snooping is enabled,

- Run `show |display set|display inheritance |match igmp-snooping`

To add the configuration after the device boots up with new Junos OS software and before enabling any interfaces,

- Run `set protocols evpn leave-sync-route-oldstyle`

Configure EVPN Type5 Enabled Devices

This section is only applicable to EVPN Type5 VRF enabled devices running Junos OS releases earlier than 19.4R1 and plan to upgrade to Junos OS 19.4R1 or later releases. In this scenario, Type-5 VRF has IPv4/IPv6 routes learnt from other BGP families. In some cases, Type-5 routes are advertised with additional BGP communities, other than EVPN communities, such as encapsulation, router-max ext community. The policies are implemented against these additional BGP communities. Before re-enabling any interfaces, you must add additional configuration after the device boots up with the new Junos OS

With the following configuration, it will not inherit communities in either [IP Prefix→EVPN] and [EVPN→IP Prefix] direction and keep the same behaviors for Junos OS releases earlier than 19.4R1.

See Junos OS EVPN User Guide for more information.

- Add the configuration after the device boot up with new Junos OS software and before enabling any interfaces:

```
set routing-instances <type-5 l3vrf instance> protocols evpn ip-prefix-routes route-attributes community export-action skip

set routing-instances <type-5 l3vrf instance> protocols evpn ip-prefix-routes route-attributes community import-action skip
```

Caveats

- This upgrade procedure expects the devices to be managed via out-of-band and via the management interface. When we shut down all ports, this action breaks management access if currently done in-band.
- Always connecting the hosts in the network to the shared ESI among multiple leafs to minimize the traffic lost during any network events.
- Single home connected host is not covered .
- Always upgrade one device at a time at initial phase. After a few successful upgrades and getting familiar with the procedure, upgrade can be done in batches with multiple single members of the leafs pairs at a time for large deployment.
- Disable HA knobs from device configuration before upgrade and enable HA knobs back after finishing the upgrade.
- Traffic lost or duplicated traffic is observed during the network upgrade procedure.
- ERB network converge time is relatedly faster compared to CRB deployment, due the local routing and switching on leaf node.
- The hosts in CRB network experience short period of traffic impact. Same behaviors on Bridge Overlay deployment too.
    - In both Spine and Leaf upgrade phase.
    - Short period of traffic lost is observed.
    - Duplicated copies of the intra-vlan traffic might be received before re-learning the host mac and ip address on upgrading devices.

Junos OS Software Rollback for EVPN Upgrade

NOTE: Only rollback Junos OS software due to upgrade failure.

- Contact customer support for alarms and core-dumps during and after the upgrade.
    - Provide the syslog file “messages” and core files.
    - Provide RSI info capture on the devices. This is applicable from Junos OS 21.2R3-S3, Junos OS 21.4R3, and later releases.
      The execution time for this command varies based on the scale of the devices.

      ```
      request support information evpn-vxlan | tee /var/tmp/rsi.txt
      ```

- Contact customer support for rolling back the Junos OS software, if upgrade fails

  ```
  request system software rollback
  request system reboot
  ```