Juniper Networks | Engineering Simplicity

IN FOCUS

Junos® OS Release 20.1

Published
2020-03-27

junos

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

IN FOCUS Junos® OS Release 20.1

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

1 **Start Here with Junos OS Release 20.1**

What You Need to Know About the In Focus Guide | 5

Important Features in Junos OS Release 20.1 | 5

2 **Segment Routing for the Path Computation Element Protocol**

How to Configure Segment Routing for the Path Computation Element Protocol | 8

Segment Routing for the Path Computation Element Protocol Overview | **8**

Benefits of Segment Routing for PCEP | **8**

Segment Routing for Traffic Engineering | **9**

Junos OS Implementation of Segment Routing for PCEP | **9**

Segment Routing for PCEP Limitations and Unsupported Features | **15**

3 **Sequential Upgrade**

How to Use Sequential Upgrade in an MX Series Virtual Chassis | 18

Sequential Upgrade Overview | **19**

Benefits of Performing a Sequential Upgrade in a MX Series Virtual Chassis | **19**

Prerequisites for Performing a Sequential Upgrade in a MX Series Virtual Chassis | **19**

Performing a Sequential Upgrade in a MX Series Virtual Chassis | **20**

How Sequential Upgrade Works in a MX Series Virtual Chassis | **21**

4 **Unified ISSU with Enhanced Mode**

How to Use Unified ISSU with Enhanced Mode | 24

Unified ISSU with Enhanced Mode Overview | **25**

Benefits of Unified ISSU with Enhanced Mode | **25**

Prerequisites for Performing Unified ISSU with Enhanced Mode | **25**

Performing Unified ISSU with Enhanced Mode | **26**

1

CHAPTER

Start Here with Junos OS Release 20.1

What You Need to Know About the In Focus Guide | 5

Important Features in Junos OS Release 20.1 | 5

What You Need to Know About the In Focus Guide

Use this guide to quickly learn about the most important features in Junos OS Release 20.1 and how you can deploy them in your network.

You might also be interested in seeing the complete list of features in the Release Notes for Junos OS Release 20.1. In addition to this guide, you can find detailed information on concepts, configuration, and examples in the Junos OS documentation.

Want to tell us what you think about this guide? E-mail us at techpubs-comments@juniper.net.

Important Features in Junos OS Release 20.1

For details on these features, go to the other chapters in this guide or click the link in the feature description below.

- **Delegate segment routing LSPs to a PCE (MX Series)**—Starting in Junos OS Release 20.1R1, you can enable a Path Computation Client (PCC) to delegate locally configured IPv4 non-colored segment routing LSPs to a Path Computation Element (PCE) controller. The PCE controls the delegated LSPs and can modify LSP attributes for traffic steering.

  A PCC with delegation capability can take back control of the delegated segment routing LSPs from the PCE when the PCEP session goes down; the LSPs would otherwise be deleted from the PCC. You can thus ensure LSP data protection by averting a situation where packets are silently discarded or dropped (also known as a traffic black-hole condition).

  [See "How to Configure Segment Routing for the Path Computation Element Protocol" on page 8.]

- **Sequential upgrade for Virtual Chassis (MX240, MX480, MX960, and MX10003)**—Starting in Junos OS Release 20.1R1, MX Series Virtual Chassis configurations can use sequential upgrade to install new software releases with minimal network downtime. Sequential upgrade is an alternative to unified ISSU that installs a new release and reboots each Virtual Chassis member router one at a time. While the upgrade installs on one member router, the other member router continues to operate and handle network operations.

  To perform a sequential upgrade in an MX Series Virtual Chassis, you first issue the **request virtual-chassis upgrade protocol-backup *package-name*** command from the CLI for the Virtual Chassis master router. This initiates the upgrade process on the Virtual Chassis backup router. After the upgrade finishes on the backup router, issue the **request virtual-chassis upgrade protocol-master *package-name*** command from the backup router CLI to begin the same upgrade process for the Virtual Chassis master router.

  [See "How to Use Sequential Upgrade in an MX Series Virtual Chassis" on page 18.]

- **Unified ISSU with enhanced mode (MX240, MX480, MX960, MX2008, MX2010, MX2020)**—Starting in Junos OS Release 20.1R1, MX Series routers with MPC7E, MPC8E, or MPC9E line cards installed can use a new ISSU option called *enhanced mode*. Enhanced mode eliminates packet loss during the unified ISSU process by running a second copy of the Junos software in standby mode ready to take over when software moves from an old image to a new one.

  Use the **request system software in-service-upgrade *package-name*.tgz enhanced-mode** command to use unified ISSU with enhanced mode, or the **request system software validate in-service-upgrade *package-name*.tgz enhanced-mode** command to verify that your device and target release are compatible with enhanced mode.

  [See "How to Use Unified ISSU with Enhanced Mode" on page 24.]

- **Software support for SRX380 devices**—The SRX380 Services Gateway is a high performance and all-in-one networking device, which consolidates routing, switching, and security. With next-generation firewall features and advanced threat mitigation capabilities, the SRX380 device provides cost-effective and secure connectivity across distributed enterprise locations. A 1U form factor model with a 16 core MIPS processor and 4 GB DDR4 RAM, the SRX380 device supports up to 10 Gbps firewall performance.

  The SRX380 device has an integrated 100-GB SSD and provides high port density with 16 on-board PoE-enabled 1-Gigabit Ethernet RJ-45 ports and 4 10-Gigabit Ethernet SFP+ ports. All the ports support AES-256 MACsec encryption. The SRX380 device has dual AC power supplies and supports up to four Mini-PIMs.

  The SRX380 supports the same features as those supported on the existing SRX300 line of services gateways. For the complete list of features supported on the SRX380, see Feature Explorer.

  [See psu and show chassis power-budget-statistics.]

2

CHAPTER

Segment Routing for the Path Computation Element Protocol

How to Configure Segment Routing for the Path Computation Element Protocol | 8

How to Configure Segment Routing for the Path Computation Element Protocol

SUMMARY

You can enable segment routing or Source Packet Routing in Networking (SPRING) traffic-engineering (SR-TE) with the Path Computation Element Protocol (PCEP) for traffic steering. With this support, the advantages of segment routing are extended to the label-switched paths (LSPs) that are externally controlled by a Path Computation Element (PCE).

IN THIS SECTION

- Segment Routing for the Path Computation Element Protocol Overview | 8

Segment Routing for the Path Computation Element Protocol Overview

IN THIS SECTION

- Benefits of Segment Routing for PCEP | 8
- Segment Routing for Traffic Engineering | 9
- Junos OS Implementation of Segment Routing for PCEP | 9
- Segment Routing for PCEP Limitations and Unsupported Features | 15

Benefits of Segment Routing for PCEP

- Setting up of LSPs through an external controller provides a global view of per-LSP and per-device bandwidth demand on the network, enabling online and real-time constraint-based path computation.

  The advantages of segment routing are extended to the LSPs initiated by an external controller, also known as a Path Computation Element (PCE), augmenting the benefits of external path computing in an MPLS network.

- A Path Computation Client (PCC, an ingress MX Series router) with delegation capability can take back control of the delegated segment routing LSPs from the PCE when the PCEP session goes down; the LSPs would otherwise be deleted from the PCC. You can thus ensure LSP data protection by averting a situation where packets are silently discarded or dropped (also known as a traffic black-hole condition).

Segment Routing for Traffic Engineering

Segment routing can operate over an IPv4 or IPv6 data plane, and supports equal-cost multipath (ECMP). With the IGP extensions built into it, segment routing integrates with the rich multiservice capabilities of MPLS, including Layer 3 VPN, Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), and Ethernet VPN (EVPN).

Some of the high-level components of the segment routing-traffic engineering (SR-TE) solution include:

- Use of an IGP for advertising link characteristics. This functionality is similar to RSVP-TE.
- Use of Constrained Shortest Path First (CSPF) on the ingress device or the PCE.
- Use of an IGP for advertising labels for links.

  In SR-TE functionality:

  1. The ingress device constructs an LSP by stacking the labels of the links that it wants to traverse.
  2. The per-link IGP advertisement is combined with label stacking to create source routed LSPs on the ingress device, so the transit devices are not aware of the end-to-end LSPs.
  3. LSPs are created between edge nodes without placing any per-LSP memory requirements on the transit devices. (Creation of such LSPs is enabled as there is no per-LSP signaling in SR-TE.)
  4. Per-neighbor labels are stacked, which results in the management of a large number of labels, leading to control plane scaling.

Junos OS Implementation of Segment Routing for PCEP

IN THIS SECTION

- PCE-Initiated Segment Routing LSPs | **10**
- PCE-Delegated Segment Routing LSPs | **11**

Junos OS implements segment routing for PCEP for two types of LSPs—PCE-initiated LSPs and PCE-delegated LSPs.

PCE-Initiated Segment Routing LSPs

The PCE-initiated segment routing LSPs are those LSPs that the PCE creates for the adjacency and node segmentsadjacency and node segments.

The PCE performs the following functions:

1. Computes the path of the segment routing LSP.

2. Provisions the LSP on the Path Computation Client (PCC) using PCEP segment routing extensions.

3. Parses the PCEP segment routing extensions.

4. Creates a tunnel route on the PCC that has its own preference value and is made available in the inet.3 routing table to resolve IP traffic and services like any other tunnel route.

The PCC performs the following functions:

1. Selects the outgoing interface based on the first network access identifier (NAI) in the source Explicit Route Object (S-ERO).

   Junos OS supports S-EROs that contain the first hop as a strict hop; Junos OS doesn't support selection of the outgoing interface on the PCC based on a loose-hop node segment ID (SID). However, the remaining hops can be loose. No specific processing is done for the S-EROs that are beyond the first hop, other than to simply use the label for next-hop creation.

2. Rejects the S-ERO if:

   - The S-ERO does not have labels in it.
   - The S-ERO carries more than six hops.

   Creates an equal-cost multipath (ECMP) route when there are multiple LSPs to the same destination with the same metric.

3. Waits for the PCE to process any event that leads to a change in the segment routing LSP after it is provisioned--for example, if the label is changed or withdrawn, or if one of the interfaces traversed by the LSP goes down.

When the PCEP session goes down, the PCE-initiated segment routing LSP:

1. Remains up for 300 seconds.

2. The LSP is deleted from the PCC after 300 seconds.

For more details, see Internet drafts draft-ietf-pce-lsp-setup-type-03.txt (expires December 25, 2015), *Conveying path setup type in PCEP messages*, and draft-ietf-pce-segment-routing-06.txt (expires February 10, 2016), *PCEP Extensions for Segment Routing*.

PCE-Delegated Segment Routing LSPs

The PCE-delegated segment routing LSPs are those LSPs that the PCC configures locally and then delegates to a PCE controller.

> **NOTE:**
> Junos OS Release 20.1R1 supports:
>
> - PCE delegation capability only for non-colored segment routing LSPs with IPv4 destinations.
> - Delegation and reporting of only the first segment of a segment list to an external controller. Multiple segments are not supported for PCE delegation.

The PCC can delegate a segment routing LSP to an external controller (the PCE) in the following ways:

- **Initial delegation**—The local LSPs are yet to be configured on the PCC, and the delegation of the LSP happens at the time the LSP is configured.
- **Delegation of existing LSP**—The local LSPs are configured on the PCC, and the delegation of the LSP happens after the source-routing path is configured. That is, the delegation capability is enabled on existing segment routing LSPs.

After delegating a segment routing LSP, the PCE controls the delegated LSPs and can modify the LSP attributes for path computation. The LSP control reverts back to the PCC when the PCEP session between the PCC and the PCE goes down. The PCE-delegated LSPs have an advantage over PCE-initiated LSPs in case the PCEP session goes down. For PCE-initiated LSPs, when the PCEP session is down, the LSPs are deleted from the PCC. However, for PCE-delegated LSPs, when the PCEP session goes down, the PCC takes back control of the delegated LSPs from the PCE. As a result, with PCE-delegated LSPs, we avert a situation where packets are silently discarded (also known as a traffic black-hole condition) when the session goes down.

The following types of segment routing LSPs support the PCE-delegation capability:

- **Static LSPs**—Statically configured source-routing paths that have the entire label stack statically configured.
- **Auto-translated LSPs**—Statically configured source-routing-paths that are automatically translated.
- **Computed LSPs**—Statically configured source-routing-paths that are computed with distributed Constrained Shortest Path First (CSPF).
- **Dynamic LSPs**—Dynamically created tunnels triggered through the Dynamic Tunnel Module that have last-hop ERO resolution.

Depending on the source of the segment routing LSP, you can configure the delegation capability on the PCC. To enable delegation of segment routing LSPs, include the **lsp-external-controller *pccd*** statement at the appropriate level under the **[edit protocols source-packet-routing]** hierarchy.

Table 1 on page 12 shows a mapping of the LSP source to the corresponding configuration hierarchy level at which the delegation capability is enabled.

> **NOTE:** You must include the **lsp-external-controller *pccd*** statement at the **[edit protocols source-packet-routing]** and **[edit protocols mpls]** hierarchy levels before configuring the delegation capability on the PCC.

Table 1: Mapping of Segment Routing LSP Source with Configuration Hierarchy

Source of Segment Routing LSP	Configuration Hierarchy
• Auto-translated LSPs • Static LSPs	Primary segment list at **[edit protocols source-packet-routing source-routing-path *lsp-name* primary *path-name*]**
Computed LSPs (distributed CSPF)	Primary segment list of the source-routing path at: • **[edit protocols source-packet-routing source-routing-path *lsp-name* primary *path-name* compute *profile-name*]** • **[edit protocols source-packet-routing source-routing-path *lsp-name* primary *path-name*]**
Dynamic LSPs	Primary segment list of the source-routing path template at: • **[edit protocols source-packet-routing source-routing-path-template *template-name* primary *primary-segment-list-name*]** • **[edit protocols source-packet-routing source-routing-path-template *template-name*]**

You can view the control status of the SR-TE LSPs from the *show spring-traffic-engineering* command output.

Table 2 on page 13 displays the PCEP interaction when the **lsp-external-controller** statement is configured for a source-routing path.

Table 2: PCEP Interaction LSP Delegation

lsp-external-controller Configuration Hierarchy	source-routing-path Delegation State	PCEP Interaction Between PCC and PCE
Primary segment list of source-routing path	Initial delegation	1. A PCReport is sent to the PCE for delegation. PCReport contains only constraints and path details (such as ERO). 2. PCE calculates the path for LSP and reports path to be in the down state. 3. No route is programmed by the local LSP until the controller computes the ERO and notifies the result to the PCC through PCUpdate. The same behavior is seen when the routing protocol process (rpd) restarts or a Routing Engine switchover happens.

Table 2: PCEP Interaction LSP Delegation *(continued)*

lsp-external-controller Configuration Hierarchy	source-routing-path Delegation State	PCEP Interaction Between PCC and PCE
Primary segment list of source-routing path	Delegation of existing path	1. A PCReport is sent to the PCE for delegation. PCReport contains only constraints and path details (such as ERO). 2. A corresponding primary segment is delegated to the PCE. 3. PCE calculates the path for the LSP. 4. The primary segment continues to contribute to the route as determined by the local configuration or computation until a PCUpdate is received from the PCE. • If Seamless BFD (S-BFD) is not configured for the primary segment, then there is no further update to the route and the LSP state is also not monitored and reported to the PCE. The LSP state at this point is reported as up or down depending on whether path computation had succeeded at that point. • If S-BFD is configured for the primary segment, then the state of the primary segment is tracked and reported to the PCE. If BFD detects the primary segment to be down, the corresponding primary path is removed from the route. The same route that was previously calculated is reprogrammed if that path is up now. 5. If a PCUpdate message is received from the PCE, SR-TE uses the received parameter to set up the path for which the PCReport message was sent. The programmed path then includes only the segment list received from the PCE and all the other segment lists that were previously programmed are removed. This reprogramming of the route happens in a make-before-break fashion.
Primary segment of source-routing path	Delegation is not configured or has been deleted	The segment list from the PCE (if available) is no longer used and the computation result from the local configuration is used. When the local result for the segment list is available, the corresponding segment list is used to program the route in a make-before-break fashion.

Table 2: PCEP Interaction LSP Delegation *(continued)*

lsp-external-controller Configuration Hierarchy	source-routing-path Delegation State	PCEP Interaction Between PCC and PCE
Segment list of source-routing path	Delegation is enabled after LSP is configured	Delegation functionality is triggered for the primary segment list under the source-routing path.
Segment list of source-routing path	Delegation is not configured or has been deleted	Delegation functionality is removed from the primary segment list under the source-routing path.
Primary segment list of source-routing path template	Delegation is enabled after LSP is configured	• Under the source-routing path template—Delegation functionality is triggered for the entire source-routing path. Template configurations can be applied only to the Dynamic Tunnel Module. • Under the primary path in the source-routing path template—Delegation functionality is triggered for that particular primary path according to the configuration.
Primary segment list of source-routing path template	Delegation is not configured or has been deleted	Delegation functionality is removed from all the source-routing paths and primary paths that match the template configuration.

Segment Routing for PCEP Limitations and Unsupported Features

The support of segment routing for PCEP does not add to the performance burden on the system; however, it has the following limitations:

- An SR-TE LSP is not locally protected on the PCC. When the LSP is more than six hops, no service is provided on the LSP other than to carry plain IP traffic.
- Graceful Routing Engine switchover (GRES) and unified in-service software upgrade (unified ISSU) are not supported.
- Nonstop active routing (NSR) is not supported.
- IPv6 is not supported.
- PCE-delegated LSPs does not support the following:
  - Colored SR-TE LSPs
  - IPv6 LSPs

- Secondary segment list of the source-routing path. Only one path of the segment list can be delegated.
- Multisegment standard. Only the first segment of the segment list is delegated and reported to the controller.

WHAT'S NEXT

> For more information on configuring Segment Routing for the Path Computation Element Protocol, see Example: Configure Segment Routing for the Path Computation Element Protocol.

3

CHAPTER

Sequential Upgrade

How to Use Sequential Upgrade in an MX Series Virtual Chassis | 18

How to Use Sequential Upgrade in an MX Series Virtual Chassis

SUMMARY

Use this document to learn about sequential upgrade, how it works, and how to initiate an sequential upgrade on MX Series Virtual Chassis configurations.

IN THIS SECTION

- Sequential Upgrade Overview | **19**
- Benefits of Performing a Sequential Upgrade in a MX Series Virtual Chassis | **19**
- Prerequisites for Performing a Sequential Upgrade in a MX Series Virtual Chassis | **19**
- Performing a Sequential Upgrade in a MX Series Virtual Chassis | **20**
- How Sequential Upgrade Works in a MX Series Virtual Chassis | **21**

Sequential Upgrade Overview

Starting in Junos OS Release 20.1R1, MX Series Virtual Chassis configurations can use sequential upgrade to install new software releases with minimal network downtime. The sequential upgrade process is an alternative to unified in-service software upgrade (ISSU) that installs a new release and reboots each Virtual Chassis member router one at a time. While the upgrade installs on one member router, the other member router continues to operate and handle network operations. This lets you upgrade to a new release with minimal disruption to the network.

Benefits of Performing a Sequential Upgrade in a MX Series Virtual Chassis

Performing a sequential upgrade in an MX Series Virtual Chassis provides the following benefits:

- Upgrades the Junos OS software package while maintaining subscriber sessions
- Minimizes network downtime during software image upgrades
- Avoids upgrading all Flexible PIC Concentrators (FPCs) and both chassis at the same time

Sequential upgrade is an alternative to unified ISSU. Compared to ISSU, sequential upgrade offers the following benefits:

- Easier troubleshooting. Sequential upgrade applies the upgrade to the backup router first, giving you a window to check on the success of the upgrade and troubleshoot if necessary
- Ability to back out of an upgrade. With sequential upgrade, you can issue the **request virtual-chassis upgrade cancel** command after the backup router is upgraded, giving you the flexibility to back out of an upgrade and roll back to the original software version
- Lower resource requirements for sequential upgrade

Prerequisites for Performing a Sequential Upgrade in a MX Series Virtual Chassis

Before you start a sequential upgrade in a two-member MX Series Virtual Chassis, make sure you do all of the following:

- Ensure that all four Routing Engines in the Virtual Chassis (both Routing Engines in the master router and both Routing Engines in the backup router) are running the same Junos OS software release.
- For minimum traffic disruption, make sure that both member routers are configured with symmetric network interface configurations so traffic can continue to run on all interfaces after switching from the master router to the backup router.
- Ensure that your network is configured to enable moving all traffic from one member router to the other.
- Back up the existing router configuration.
- Verify that both graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are enabled.
- Disable load throttling with the **set system services resource-monitor no-load-throttle** command. This will allow subscriber logins during the upgrade process.
- Download the software release package you want to upgrade to from Juniper's website at https://support.juniper.net/support/downloads/

> **NOTE:** We do not recommend manually moving the active interface to the backup router. Instead allow the configured LACP failure detection and switchover mechanism to do this when the master router is rebooted during its upgrade procedure.

Performing a Sequential Upgrade in a MX Series Virtual Chassis

To perform a sequential upgrade in an MX Series Virtual Chassis, follow these steps:

1. From the Virtual Chassis master router, issue the following CLI command to initiate the upgrade process on the Virtual Chassis backup router:

   request virtual-chassis upgrade protocol-backup *package-name*

   > **NOTE:** Between upgrading the backup router and master router, the statistics for the backup router are not properly displayed on the master router CLI. Statistics for the backup router are available on the backup router CLI.

2. After the upgrade finishes on the backup router, issue the following CLI command from the backup router to begin the same upgrade process for the Virtual Chassis master router:

request virtual-chassis upgrade protocol-master *package-name*

> **NOTE:** You can override the automatic rebooting of Routing Engines by adding the **halt-re** statement to the CLI command as follows: **request virtual-chassis upgrade protocol-backup halt-re *package-name***. This causes the Routine Engines to halt and wait for the operator to reboot them. Use this command if you want to replace hardware components like FPCs while the router is shut down during the upgrade process.

If you want to cancel the sequential upgrade process, you can issue the **request virtual-chassis upgrade cancel** command on any Routing Engine after the Virtual Chassis backup router is updated and before the Virtual Chassis master router is updated. After canceling the upgrade process, use the **request system software rollback** command to rollback to the previously installed package, if necessary.

How Sequential Upgrade Works in a MX Series Virtual Chassis

At a high level, the software performs the following actions after you issue the **request virtual-chassis upgrade protocol-backup *package-name*** command to upgrade to a new Junos OS software release in a two-member Virtual Chassis configuration:

1. Verifies that an upgrade is not already in progress and exits if an upgrade is already running.
2. Validates the software image using existing installation support.
3. Performs configuration validation.
4. Copies the software image to the backup router Routing Engines.
5. Activates the software image on the backup router Routing Engines.
6. Reboots the backup router Routing Engines.
7. Polls the backup router for FPC interface synchronization after the Routing Engines reboot.

After upgrading the Virtual Chassis backup router, the next step is to issue the **request virtual-chassis upgrade protocol-master *package-name*** command from the backup router CLI to initiate the upgrade process on Virtual Chassis master router. The software performs the same actions as listed in steps 1-7 above, but for the master router Routing Engines. After the master router Routing Engines reboot, the software performs the following actions:

- Switches the role of the Virtual Chassis backup router to become the Virtual Chassis master router.

WHAT'S NEXT

> Now that you've learned about performing a sequential upgrade in a MX Series Virtual Chassis, you should read up on the unified ISSU process for MX Series Virtual Chassis to decide which upgrade method works best for you. Check out Upgrading Junos OS in an MX Series Virtual Chassis by Performing a Unified ISSU for more information.

Release History Table

Release	Description
20.1	Starting in Junos OS Release 20.1R1, MX Series Virtual Chassis configurations can use sequential upgrade to install new software releases with minimal network downtime.

4

CHAPTER

Unified ISSU with Enhanced Mode

How to Use Unified ISSU with Enhanced Mode | 24

How to Use Unified ISSU with Enhanced Mode

SUMMARY

Use this document to learn about unified ISSU with enhanced mode and how to use it.

IN THIS SECTION

- Unified ISSU with Enhanced Mode Overview | **25**
- Benefits of Unified ISSU with Enhanced Mode | **25**
- Prerequisites for Performing Unified ISSU with Enhanced Mode | **25**
- Performing Unified ISSU with Enhanced Mode | **26**

Unified ISSU with Enhanced Mode Overview

Enhanced mode is an in-service software upgrade (ISSU) option available on MPC7E, MPC8E, and MPC9E line cards that eliminates packet loss during the unified ISSU process. New architecture improvements on these line cards make it possible to have a second copy of the Junos OS software running on the line card in standby mode ready to take over while software moves from an old image to a new one during unified ISSU. You can enable enhanced mode by adding the **enhanced-mode** option to the **request system software in-service-upgrade** CLI command.

Benefits of Unified ISSU with Enhanced Mode

Unified ISSU with enhanced mode provides the following benefits:

- Upgrades to a new software version with no loss of transit or host bound traffic
- Reduces packet loss to zero or several milliseconds depending on configuration and network conditions
- Allows software upgrades to be performed without the need for maintenance windows
- Uses the existing unified ISSU process and doesn't require any special configuration

Prerequisites for Performing Unified ISSU with Enhanced Mode

Before you begin a unified ISSU with enhanced mode, there are several prerequisites to keep in mind:

- The device running unified ISSU with enhanced mode must use an MPC7E, MPC8E, or MPC9E line card.

  NOTE: If you are performing unified ISSU with enhanced mode on a device that has a mix of supported and unsupported line cards, there will be sub-second traffic loss for traffic passing through the unsupported line cards.

  NOTE: If you are performing unified ISSU with enhanced mode on guest network functions (GNFs), then all GNFs should be using MPC7E, MPC8E, or MPC9E line cards to avoid traffic loss.

- The Linux version running on your Flexible PIC Concentrator (FPC) and the line card Linux version in the target release need to be compatible. Use the **request system software validate in-service-upgrade *package-name.tgz* enhanced-mode** CLI command to check compatibility.
- Forwarding memory usage should be below 75 percent to ensure no packet loss during the unified ISSU process

  NOTE: Unified ISSU with enhanced mode will still work if forwarding memory usage is above 75 percent, but it might introduce several milliseconds of packet loss.

- All prerequisites for unified ISSU also apply to enhanced mode. See Unified ISSU System Requirements for more information.

You can check to see if your device can use unified ISSU with enhanced mode to upgrade to a specific release by using the **request system software validate in-service-upgrade *package-name*.tgz enhanced-mode** command. If your device and the target release are not compatible with enhanced mode, you can still use regular unified ISSU to upgrade with minimal disruption of traffic.

Performing Unified ISSU with Enhanced Mode

To perform a unified ISSU with enhanced mode, follow these steps:

1. Download the software package by following the procedure in Downloading Software.
2. Copy the software package or packages to the device. We recommend that you copy the file to the **/var/tmp** directory.

3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Verify that you can use unified ISSU with enhanced mode for your desired release.

   On the device, enter:

   ```
   user@host> request system software validate in-service-upgrade
   /var/tmp/package-name.tgz enhanced-mode
   ```

   where ***package-name*.tgz** is the name of the software package you downloaded in Step 1.
5. Start the unified ISSU with enhanced mode:

   On the device, enter:

   ```
   user@host> request system software in-service-upgrade /var/tmp/package-name.tgz
    enhanced-mode reboot
   ```

   where ***package-name*.tgz** is the name of the software package you downloaded in Step 1.

   NOTE: During the upgrade, you will not be able to access the Junos OS CLI.

   The device displays status messages similar to the following messages as the upgrade proceeds:

   ```
   Chassis ISSU enhanced-mode
   ISSU: set chassis enhanced-mode
   Chassis ISSU Check Done
   ISSU: Validating Image
   ..
   mgd: commit complete
   Validation succeeded
    Validating Image Done
    Preparing Backup RE
    Pushing /var/tmp/junos-install-mx-x86-32-20.1.tgz to
   re1:/var/tmp/junos-install-mx-x86-32-20.1.tgz
    Pushing package /var/tmp/junos-install-mx-x86-32-20.1.tgz to re1 done
    Installing package /var/tmp/junos-install-mx-x86-32-20.1.tgz on re1
   …
   Verified sflow-mx signed by PackageDevelopmentEc_2019 method ECDSA256+SHA256
   NOTICE: 'pending' set will be activated at next reboot...
   ISSU: Installing package /var/tmp/junos-install-mx-x86-32-20.1.tgz on re1 done
   ISSU: Rebooting Backup RE
   ```

```
Rebooting re1
 Backup RE Prepare Done
 Waiting for Backup RE reboot
 Backup RE reboot done. Backup RE is up
 Waiting for Backup RE state synchronization
 Backup RE state synchronization done
 GRES operational
 "Initiating Chassis In-Service-Upgrade"
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Offline Incompatible FRUs
ISSU: Starting Upgrade for FRUs
…

ISSU: Preparing for Switchover
ISSU: Ready for Switchover
 Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 1          Online (ISSU)
  FPC 2          Offline                 Configured power off
Resolving mastership...
Complete. The other routing engine becomes the master.
```

6. Log in after the reboot of the device is completed. To verify that the software has been upgraded, enter the following command:

```
user@host> show version
```

NOTE: When using unified ISSU with enhanced mode, the base Linux OS on your FPC cannot be upgraded as part of the ISSU process. Linux can be updated with an upgrade done through regular unified ISSU or a reboot of the FPC.

WHAT'S NEXT

Now that you understand how to run a unified ISSU with enhanced mode, you may want to read up on the unified ISSU process to understand what your device goes through during a unified ISSU. Check out Understanding the Unified ISSU Process for more information.