Release Notes

Published
2025-03-28

Junos OS Evolved Release 22.2R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 22.2R1.

For more information on this release of Junos OS Evolved, see Introducing Junos OS Evolved.

Table of Contents

Key Features in Junos OS Evolved Release 22.2 | 1

Junos OS Evolved Release Notes for ACX7100-32C, ACX7100-48L, and ACX7509 Devices

What's New | 5

What's New in 22.2R1 | **5**

Chassis | **6**

Dynamic Host Configuration Protocol | **6**

EVPN | **6**

Interfaces | **7**

Junos Telemetry Interface | **7**

Multicast | **8**

OpenConfig | **9**

Routing Protocols | **10**

Additional Features | **10**

What's Changed | 13

What's Changed in Release 22.2R1-S1 | **14**

What's Changed in Release 22.2R1 | **14**

Known Limitations | 19

Open Issues | 22

Resolved Issues | 25

Resolved Issues: 22.2R1 | **26**

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Devices

What's New | 29

What's New in 22.2R1 | **30**

Chassis | **31**

Class of Service | **31**

EVPN | **32**

Interfaces | **32**

IP Tunneling | 33
Juniper Extension Toolkit (JET) | 34
Junos Telemetry Interface | 34
MPLS | 37
Multicast | 37
Network Management and Monitoring | 38
OpenConfig | 41
Routing Protocols | 43
Software Installation and Upgrade | 45
Source Packet Routing in Networking (SPRING) or Segment Routing | 45
VLANs | 46
Additional Features | 46

What's Changed | 49

What's Changed in Release 22.2R1-S1 | 49

What's Changed in Release 22.2R1 | 50

Known Limitations | 53

Open Issues | 54

Resolved Issues | 58

Resolved Issues: 22.2R1 | 59

Junos OS Evolved Release Notes for QFX5130-32CD, QFX5220, and QFX5700 Devices

What's New | 65

What's New in 22.2R1 | 65

Class of Service | 65
EVPN | 66
Juniper Extension Toolkit (JET) | 67
Junos Telemetry Interface | 68
MPLS | 68
OpenConfig | 69
Routing Protocols | 69
Additional Features | 70

What's Changed | 71

What's Changed in Release 22.2R1-S1 | 72

What's Changed in Release 22.2R1 | 72

Known Limitations | 77

Open Issues | 78

Resolved Issues | 79

Resolved Issues: 22.2R1 | 80

Upgrade Your Junos OS Evolved Software | 81

Licensing | 82

Finding More Information | 82

Requesting Technical Support | 83

Revision History | 84

Key Features in Junos OS Evolved Release 22.2

Start here to learn about the key features in Junos OS Evolved Release 22.2. For more information about a feature, click the link in the feature description.

- **Blocking asymmetric EVPN Type 5 routes (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can configure the local node to reject asymmetric EVPN Type 5 routes on EVPN-VXLAN networks. The local node examines the incoming EVPN Type 5 route packets and rejects the route when the virtual network identifier (VNI) in the ingress route differs from the locally configured VNI.

  To block asymmetric EVPN Type 5 routes, include the `reject-asymmetric-vni` statement at the `[edit routing-instance` *`routing-instance-name`* `protocols evpn ip-prefix-routes]` hierarchy level.

  [See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN and ip-prefix-routes.]

- **Collect ON_CHANGE BGP RIB telemetry statistics** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5220)

  [See Telemetry Sensor Explorer.]

- **EVPN-MPLS E-LAN flow-aware transport (FAT) label load balancing (PTX10001-36MR, PTX10004, PTX10008, PTX10016)** —Starting in Junos OS Evolved 22.2R1, you can configure provider edge (PE) devices to use FAT labels in an EVPN-MPLS routing instance, according to RFC 6391. Provider edge devices use these labels to load-balance EVPN-MPLS unicast packets across ECMP paths without needing to do deep packet inspection of the MPLS payload. This feature supports E-LAN with single-homing and multi-homing active/standby and active/active topologies and supports the VLAN-based, VLAN-bundle, and VLAN-aware bundle EVPN-MPLS variants.

  To enable load balancing using FAT labels in an `evpn` routing instance:

  - Configure the `flow-label-static` statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn` hierarchy level on PE devices to insert FAT flow labels into pseudowire packets sent to remote PE devices.

  - Configure the `flow-label` statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn` hierarchy level on PE devices to signal flow-label capability in the EVPN Layer 2 Attributes Extended Community by setting the flow-label (F) bit in the EVPN Type 3 route.

  [See flow-label and flow-label-static.]

- **gRIBI sensor support (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI streams programmable routing protocol process (prpd) statistics and gRPC routing information base (RIB, also known as routing table) programming interface-related statistics.

[See Telemetry Sensor Explorer.]

- **Interconnect EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved 22.2R1, we support Data Center Interconnect (DCI) stitching for EVPN-VXLAN gateway tunnels. The gateway connects the data center and the WAN, and both data center and WAN gain forwarding states for route distinguishing, route targeting, and interconnect Ethernet segment identifier (I-ESI) support.

  DCI control plane stitching also:

  - Supports multihoming.
  - Extends the Layer 2 (L2) connectivity required for some tenants in a data center.
  - Uses the unknown MAC route to prevent MAC scale issues on data center network virtualization edge (NVE) devices.

  [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.]

- **Queue-depth monitoring support for virtual output queues (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Virtual output queue (VOQ) queue-depth monitoring, or latency monitoring, measures peak queue occupancy of a VOQ. Starting with Junos OS Evolved Release 22.2R1, PTX Series routers running Junos OS Evolved support VOQ queue-depth monitoring to report peak queue length for a given physical interface for each individual Packet Forwarding Engine.

  [See VOQ Queue-depth Monitoring.]

- **Support for CoS within MPLS networks (QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use CoS within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. We have also added MPLS EXP rewrite support.

  - **Default CoS on the provider (P) and provider edge (PE) routers for MPLS interfaces**–The MPLS traffic uses the default EXP classifier. MPLS traffic is treated as best-effort traffic using the 802.1 default untrusted classifier. The default EXP classifier applies to all MPLS traffic on interfaces configured as `family mpls`. Differentiated Services code point (DSCP) classifiers are not applied to MPLS traffic.
  - **Default CoS on PE routers for Layer 3 interfaces**–By default, all Layer 3 VPN logical interfaces are bound to default DSCP classifiers.

  If you apply an EXP classifier on a penultimate-hop popping (PHP) node, then by default, the MPLS header TLL value overwrites the IP header time-to-live (TTL) value. In this case, a zero (0) overwrites the IP header DSCP bits, which signifies uniform mode. To use pipe mode, where nothing overwrites IP header TTL values or IP header DSCP bits, you should configure the following command:

```
set protocols mpls no-propagate-ttl
```

NOTE: The DSCP of IP in MPLS packets can't be remarked either at PE or P routers.

[See Understanding CoS MPLS EXP Classifiers and Rewrite Rules.]

- **Support for DHCPv6 on zero-touch provisioning (ZTP) (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, ZTP supports the DHCPv6 client on management and WAN interfaces. During the bootstrap process, the device first uses the DHCPv4 client to request information regarding image and configuration files from the DHCP server. The device checks the DHCPv4 bindings sequentially. If one of the DHCPv4 bindings fails, the device continues to check for bindings until provisioning is successful. If no DHCPv4 bindings exist, however, the device checks for DHCPv6 bindings and follows the same process as for DHCPv4 until the device can be provisioned successfully. The DHCP server uses DHCPv6 options 59 and 17 and applicable sub-options to exchange ZTP-related information between itself and the DHCP client.

  See [Zero Touch Provisioning].

- **Support for EVPN-virtual private wire service (VPWS) (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—We've extended support for EVPN-VPWS to the listed PTX Series routers. By default, control-word is enabled on these platforms. To disable the control word feature, use the `set routing-instances` *`routing-instance-name`* `protocols evpn no-control-word` command.

  [See Overview of VPWS with EVPN Signaling Mechanisms.]

- **Support for the gRPC Network Operations Interface (gNOI)** `CertificateManagement` **service (PTX10008 and PTX10016)**—Starting in Release 22.2R1, Junos OS Evolved devices support gRPC Network Operations Interface services. You can execute supported CertificateManagement service remote procedure calls (RPCs) to manage certificates on the network device. Using gNOI operations enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

  [See gNOI Certificate Management Service.]

- **Symmetric IRB with EVPN Type 2 routes (ACX7100, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can enable symmetric IRB EVPN Type 2 routing in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) edge-routed bridging (ERB) overlay fabric. With the symmetric routing model, leaf devices can route and bridge traffic on both ingress and egress sides of a VXLAN tunnel. Leaf devices use a transit VXLAN network identifier (VNI) and Layer 3 (L3) interfaces on the associated VLAN to exchange traffic across the VXLAN tunnels.

We support this feature with `vlan-aware` and `vlan-based` MAC-VRF instance service type configurations. To enable this feature, you must also configure EVPN Type 5 routing with L3 VRF instances to establish intersubnet reachability among the EVPN devices.

[See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics and irb-symmetric-routing.]

- **TCP authentication option (TCP-AO) for resource public key infrastructure (RPKI) validation sessions (ACX7100-32C, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use the TCP authentication option to authenticate RPKI validation sessions for securing the Internet's routing infrastructure, such as the BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

  To enable a TCP authentication option chain to authenticate an RPKI validation session, use the configured `authentication-algorithm` **`ao`** `and authentication-key-chain` *`keychain`* at the [`edit routing-options validation group` *`group_name`* `session` *`address`*] and [`edit routing-options validation group` *`group_name`*] hierarchy level.

  [See TCP Authentication Option (TCP-AO)]

Junos OS Evolved Release Notes for ACX7100-32C, ACX7100-48L, and ACX7509 Devices

IN THIS SECTION

- What's New | **5**
- What's Changed | **13**
- Known Limitations | **19**
- Open Issues | **22**
- Resolved Issues | **25**

These release notes accompany Junos OS Evolved Release 22.2R1 for ACX7100-32C, ACX7100-48L, and ACX7509 devices. The release notes describe new features, limitations, and known problems in the hardware and software.

What's New

IN THIS SECTION

- What's New in 22.2R1 | 5

Learn about new features introduced in this release for ACX Series routers.

What's New in 22.2R1

IN THIS SECTION

- Chassis | 6
- Dynamic Host Configuration Protocol | 6
- EVPN | 6
- Interfaces | 7
- Junos Telemetry Interface | 7
- Multicast | 8
- OpenConfig | 9
- Routing Protocols | 10
- Additional Features | 10

To view features supported on the ACX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 22.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- ACX7100-32C
- ACX7100-48L
- ACX7509

The following sections highlight the key features in this release.

Chassis

- **High Availability (HA) Platform Redundancy Support for ACX7509 Devices**— High availability platform redundancy support for ACX7509 device includes the following:
  - Support for Graceful RE Switchover (GRES) for simultaneous switchover of RCB and FEB. Support for nonstop active routing (NSR) seamless switchover with Layer-2 or Layer-3 protocol data packets.
  - Support for two RCBs and two FEBs for redundancy. The RCB and FEB mastership is tied. The RCB0-FEB0 and RCB1-FEB1 switchover together. RCB and FEB are a host-subsystem and switchover due to one will cause a complete host sub-system switch.
  - ACX7509 has one CPU per RE, and transmit traffic (including transit traffic, host-traffic, inline keepalive sessions) is affected when the master RE or FEB is affected. The recovery action is a fast RE switchover with support for replication and restoration of statistics.

  [See No Link Title,No Link Title,No Link Title,No Link Title,No Link Title,No Link Title, and No Link Title.]

Dynamic Host Configuration Protocol

- **DHCPv4 and DHCPv6 relay agent support for MC-LAG (ACX7100-32C and ACX7100-48L)**— Starting in Junos OS Evolved Release 22.2R1, DHCPv4 and DHCPv6 relay agents support multichassis link aggregation groups (MC-LAG). These DHCP relay agents are used to forward DHCP requests and replies between clients and servers when they are on different physical subnets. In this release, we support the stateless `forward-only` DHCP relay agent mode.

  DHCP relay agent support includes the following:

  - DHCPv4 and DHCPv6 Stateless `forward-only` option on Layer 3 static interfaces over MC-LAG
  - DHCPv4 and DHCPv6 Stateless `forward-only` option on IRB interfaces over MC-LAG
  - DHCPv4 and DHCPv6 `forward-snooped-clients` on dual-stack configurations

  [See DHCP Relay Agent and Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent.]

EVPN

- **Support for overlapping VLANs on the same leaf device (ACX7100-32C and ACX7100-48L)**— Starting in Junos OS Evolved Release 22.2R1, you can configure overlapping VLANs in enterprise and service provider style CLIs. In cloud-based Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS) deployments, VLANs can overlap between the private cloud network and customer traffic. For example, the private cloud network can use a set of VLANs connected to a physical port

on a leaf device. You can then connect a customer VLAN on different physical ports on the same leaf device. In this scenario, you need to forward the customer traffic over the same VXLAN core network.

[See Ethernet Switching Guide].

- **Symmetric IRB with EVPN Type 2 routes (ACX7100, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can enable symmetric IRB EVPN Type 2 routing in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) edge-routed bridging (ERB) overlay fabric. With the symmetric routing model, leaf devices can route and bridge traffic on both ingress and egress sides of a VXLAN tunnel. Leaf devices use a transit VXLAN network identifier (VNI) and Layer 3 (L3) interfaces on the associated VLAN to exchange traffic across the VXLAN tunnels.

  We support this feature with `vlan-aware` and `vlan-based` MAC-VRF instance service type configurations. To enable this feature, you must also configure EVPN Type 5 routing with L3 VRF instances to establish intersubnet reachability among the EVPN devices.

  [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics and irb-symmetric-routing.]

Interfaces

- **Support for Layer 2 bridging and Layer 3 unicast (ACX7509)**—Starting in Junos OS Evolved Release 22.2R1, the following features are available on multichassis link aggregation groups (MC-LAGs):

  - Layer 2 bridging for active-active and active-standby modes
  - Layer 3 unicast

  [See Understanding Multichassis Link Aggregation Groups.]

- **Support for 50GbE interfaces (ACX7100-48L)**—Starting in Junos OS Evolved Release 22.2R1, we support 50GbE interfaces.

  [See Port Speed on ACX7100-48L Router Overview.]

Junos Telemetry Interface

- **Interface queue statistics sensor support (ACX7100, ACX7509)—**Starting in Junos OS Evolved Release 22.2R1, Juniper telemetry interface (JTI) provides periodic streaming and ON_CHANGE support for interface queue statistics. The resource path is /junos/system/linecard/interface/queue/. This feature uses Juniper proprietary gRPC and gNMI.

  [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

- **Network instance support enhancements (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.2R1, JTI supports new sensors for network instance statistics per OpenConfig modules openconfig-network-instance.yang and openconfig-routing-policy.yang. Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **Two-Way Active Measurement Protocol sensor support (ACX7100-32C, ACX7100-48L PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, Junos telemetry interface (JTI) provides periodic streaming support for Two-Way Active Measurement Protocol (TWAMP) statistics. This feature supports IPv4 and IPv6 configurations for TWAMP managed sessions and TWAMP light sessions. This feature is supported using Juniper proprietary gRPC and gNMI.

  [See Telemetry Sensor Explorer.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

- **Xmlproxyd process now multi-threading capable (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, the xmlproxyd process is now multi-threading capable. Moving from a single-thread capability to a multi-thread capability enhances the overall throughput of xmlproxyd and makes xmlproxyd scalable. Performance bottlenecks are reduced since better throughput improves streaming performance.

  [See Configure a NETCONF Proxy Telemetry Sensor in Junos.]

Multicast

- **Support for multichassis link aggregation groups (MC-LAGs) (ACX7100-32C and ACX7100-48L)**—Starting in Junos OS Evolved Release 22.2R1, the following features are available on MC-LAGs:

  - Layer 2 bridging for active-active and active-standby modes

- Layer 2 multicast with and without IGMP or MLD snooping
- Layer 3 multicast with and without IGMP or MLD snooping

[See No Link Title.]

OpenConfig

- **Network instance support enhancements (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**—Junos OS Evolved Release 22.2R1 supports new sensors for network instance statistics based on OpenConfig (OC) modules **openconfig-network-instance.yang** and **openconfig-routing-policy.yang**. Support includes OC configuration and streaming of operational state data.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **OpenConfig MPLS label distribution protocol (LDP) configuration support (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)—**—Junos OS Evolved Release 22.1R1 introduces OpenConfig (OC) support for LDP configuration based on the OpenConfig data model openconfig-mpls-ldp.yang.

  [See Mapping OpenConfig MPLS Commands to Junos Configuration.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

- You configure the JSD port with the command `set services extension-service request-response grpc clear-text port` *`port-number`* `max-connections` *`number`*. You use OpenConfig to configure that port as a p4rt port. You then create a Layer 2 and Layer 3 interfaces configuration on the device with VLAN tags and the IP address for the device, allowing Layer 2 packets to be punted and injected between the device and the SDN controller. You use firewall filters to direct packet punts to the p4-switch.

  [See OpenConfig User Guide.]

Routing Protocols

- **TCP authentication option (TCP-AO) for resource public key infrastructure (RPKI) validation sessions (ACX7100-32C, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use the TCP authentication option to authenticate RPKI validation sessions for securing the Internet's routing infrastructure, such as the BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

  To enable a TCP authentication option chain to authenticate an RPKI validation session, use the configured `authentication-algorithm` **`ao`** `and authentication-key-chain` *`keychain`* at the [`edit routing-options validation group` *`group_name`* `session` *`address`*] and [`edit routing-options validation group` *`group_name`*] hierarchy level.

  [See TCP Authentication Option (TCP-AO)]

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **Accept BGP routes with accept-own community** (ACX7100-32C, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See BGP accept-own Community and accept-own.]

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding** (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)

  [See Telemetry Sensor Explorer.]

- **EVPN E-LAN Active-Active Multihoming with EVPN Aliasing Support for ESI LAG** (ACX7509)
  [See EVPN Multihoming Overview, Example: Configuring EVPN Active-Active Multihoming, Using the evpn-aliasing-optimize command, and EVPN Features in EVPNs using EVPN LAGs]

- **Firewall filter actions** (ACX7100-32C, ACX7100-48L, and ACX7509)—We support the following:

  - next-ip firewall filter action for IPv4 address
  - next-ip6 firewall filter action for IPv6 address
  - next-interface firewall filter action for an interface

  [See Filter-Based Forwarding Overview and Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address.]

- **Flow-Aware Transport for Pseudowires (FAT) label and Entropy label support for Layer 2 circuit and Layer 2 VPN**(ACX7100-32C, ACX7100-48L, and ACX7509)—

  FAT Label support includes these features:

  - Transmit and receive direction on Layer 2 circuits and Layer 2 VPNs

  - Transmit and receive direction on Layer 2 VPN
  - Layer 2 VPN with multihoming
  - Layer 2 circuits with redundant pseudowire
  - Load balancing over LAG
  - Static flow label for Layer 2 circuits

  Entropy Label support includes these features:

  - Entropy label indication for LDP and RSVP transport
  - Load balancing over LAG

  [See FAT Flow Labels Overview and Configuring Entropy Labels.]

- **GRES support for Connectivity Fault Management (CFM) and Link Fault Management (LFM)** (ACX7509)

  [See Introduction to OAM Connectivity Fault Management (CFM) and Introduction to OAM Link Fault Management (LFM).]

- **GRES support for ERPS** (ACX7509)

  [See Understanding Ethernet Ring Protection Switching Functionality.]

- **GRES support for MPLS, L3VPN, L3 Multicast, and BGP** (ACX7509)

  [See MPLS Overview.]

- **GRES support for VRRP** (ACX7509)

  [See Understanding VRRP.]

- **High availability support for BFD (centralized mode), L3 protocols (ISIS,OSPF, ECMP), static routes, SR-ISIS, and SR-OSPF** (ACX7509)

  [See Configuring BFD.]

- **High availability support for Periodic Packet Management** (ACX7509)

[See Periodic Packet Management.]

- **LACP redundancy support** (ACX7509)

  [See Configuring Link Aggregation Control Protocol.]

- **Redundancy support for Spanning Tree Protocol (STP)** (ACX7509)

  [See Spanning-Tree Protocol Overview.]

- **sFlow ingress support**(ACX7100-32C, ACX7100-48L, and ACX7509)

  [See sFlow Monitoring Technology and Understanding How to Use sFlow Technology for Network Monitoring.]

- **SPRING support for SR-TE** (ACX7509)

  - Segment routing policy to steer labeled or IP traffic at ingress routers.
  - Segment routing paths for a non-colored static label-switched path (LSP).
  - Color-based traffic steering of Layer 2 and Layer 3 VPN services.
  - SR policy-based dynamic tunnel module triggered SR-TE.
  - Indirect next hop and composite next hop modes.

  [See Segment Routing Traffic Engineering at BGP Ingress Peer Overview.]

- **Support for all-active multi-homing redundancy in both Ethernet VPN–virtual private wire service (EVPN-VPWS) and EVPN-VPWS with flexible cross-connect. Support for EVPN VPWS Multi-Homing All-Active for Segment Routing over MPLS** (ACX7509)

  [See Overview of Flexible Cross-Connect Support on VPWS with EVPN and erview of VPWS with EVPN Signaling Mechanisms.]

- **Support for G.8275.1 telecom profile, Precision Time Protocol (PTP) over Ethernet encapsulation, and hybrid mode** (ACX7100-32C|ACX7100-48L)

  [See G.8275.1 Telecom Profile, Guidelines to Configure PTP over Ethernet, and Hybrid Mode Overview .]

- **Support for Precision Time Protocol (PTP) G.8275.1 enh and G.8275.2 enh profiles with PTP over IPv4 and IPv6 unicast** (ACX7100-32C and ACX7100-48L) —The profiles also support ordinary and boundary clocks. The G.8275.1 enh profile does not support unicast negotiation.

  [See PTP profiles.]

- **Support for Layer 2 Features - MAC move limit and multiple trunk ports** (ACX7100-32C, ACX7100-48L, and ACX7509)—

The Layer 2 feature support includes:

- The MAC move limit on Layer 2 bridge domains, VXLAN, virtual private LAN service (VPLS), and EVPN networks. Juniper Networks ACX7509 routers do not support MAC move limit over VXLAN routing instances or data center logical interfaces.
- Multiple trunk ports. Juniper Networks ACX7509 routers do not support multiple trunk ports on data center logical interfaces.

  [See Understanding MAC Limiting and MAC Move Limiting.]

- **Support for multiprotocol BGP MVPN (also known as "Next Generation MVPN" or "NG-MVPN") with MDLP as Provider Tunnel. NSR and GRES support for NG-MVPN** (ACX7509)

  [See Multiprotocol BGP MVPNs Overview]

- **Support for Redundant Pseudowire Layer 2 Circuits (ACX7100-32C and ACX7100-48L). Multichassis Link Aggregation Group (MC-LAG) is not supported with redundant pseudowire Layer 2 circuits** (ACX7509)

  [See Redundant Pseudowires for Layer 2 Circuits and VPLS]

- **Supported transceivers, optical interfaces, and direct attach copper (DAC) cables** —Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update this tool and provide the first supported release information when the optic becomes available.

What's Changed

IN THIS SECTION

- What's Changed in Release 22.2R1-S1 | 14
- What's Changed in Release 22.2R1 | 14

Learn about what changed in these releases for ACX Series routers.

What's Changed in Release 22.2R1-S1

IN THIS SECTION

- Juniper Extension Toolkit (JET) | 14
- Junos OS API and Scripting | 14

Juniper Extension Toolkit (JET)

- **Python 3 is the default and only Python version for executing Juniper Extension Toolkit Python scripts (ACX Series, PTX Series, and QFX Series)**—Junos OS Evolved supports only Python 3 for executing Juniper Extension Toolkit (JET) scripts written in Python. Python 2.7 is no longer supported for executing JET scripts, and we've deprecated the `language python` statement at the `[edit system scripts]` hierarchy level.

  [See Understanding Python Automation Scripts for Junos Devices.]

Junos OS API and Scripting

- **Deprecated functions in the `libpyvrf` Python module (ACX Series, PTX Series, and QFX Series)**—The `libpyvrf`Python module no longer supports the `get_task_vrf()` and `set_task_vrf()` functions.

  [See How to Specify the Routing Instance in Python 3 Applications on Devices Running Junos OS Evolved.]

What's Changed in Release 22.2R1

IN THIS SECTION

- General Routing | 15
- Interfaces and Chassis | 15
- Layer 3 VPN | 15
- OpenConfig | 16
- Multicast | 16
- MPLS | 16
- Network Management and Monitoring | 16

- Public Key Infrastructure | 17
- Routing Protocols | 17
- Timing and Synchronization | 18
- User Interface and Configuration | 18

General Routing

- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series and QFX Series)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route `set routing-options?`. Any instance you explicitly configure is translated into `set routing-instance r1 routing-options?`. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.

- Stateful port configuration for PTP over Ethernet and default profile is supported only on boundary clock mode and not on ordinary clock mode.

- **Change in in unnumbered-address support for GRE tunnel**—There is a behavioural change in unnumbered-address support for GRE tunnel with IPv6 family and display donor interface for both IPv4 and IPv6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level.

  [See show interfaces.]

Interfaces and Chassis

- **Display the donor details of the IPv6 borrower interface**— The output for the show interfaces command now displays the donor details of the IPv6 borrower interface.

  [See show interfaces.]

Layer 3 VPN

- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series and QFX Series)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route `set routing-options ?`. Any instance you explicitly configure is translated into `set routing-instance r1 routing-options ?`. The issue appears in translation, when you

change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.

OpenConfig

- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".

Multicast

- **Changes to `show mvpn c-multicast` and `show mvpn instance` outputs**— The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

  [See show mvpn c-multicast.]

  .]

MPLS

- When defining a constrained path LSP using more than one strict hop belonging to the egress node, the first strict hop must be set to match the IP address assigned to the egress node on the interface that receives the RSVP Path message. If the incoming RSVP Path message arrives on an interface with a different IP address the LSP is rejected.

Network Management and Monitoring

- **Changes to the NETCONF `<edit-config>` RPC response (ACX Series, PTX Series, and QFX Series)**— When the `<edit-config>` operation returns an error, the NETCONF server does not emit a `<load-error-count>` element in the RPC response. In earlier releases, the `<edit-config>` RPC response includes the `<load-error-count>` element when the operation fails.

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**— The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

  [See max-datasize.]

- **Change in behavior of SNMP MIB object ifAlias**— SNMP MIB object ifAlias now shows the configured interface alias. In earlier releases, ifAlias used to show configured interface description.

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

  See privacy-3des and privacy-aes128

Public Key Infrastructure

- **Support for PKI daemon (ACX7024)**— We've introduced support for PKI (public key infrastructure) daemon on ACX7024 routers. PKI daemon provides support to create two kinds of digital certificates on the device, local certificate and CA certificate. Using local certificate, any remote resource can verify the device's identity. Using CA certificate, the device can verify the identity of remote resource by verifying the remote resource's digital certificate. A secure channel between the device and the remote resource is created only after this mutual verification. You can use the PKI configuration and operation commands to perform PKI operations on ACX7024 routers.

  [See Public Key Infrastructure (PKI), PKI Overview, and PKI in Junos OS.]

Routing Protocols

- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the `[edit system services ssh]` hierarchy level.

  In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the `[edit system services ssh]` hierarchy level.

  [See services (System Services).]

- The RPD_OSPF_LDP_SYNC message not logged?On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the `hold-time` for `ldp-synchronization` at the **edit protocols ospf area area id interface interface name** hierarchy level less than three minutes. The message is printed after three minutes.

- To achieve consistency among resource paths, the resource path/**/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter ip-addr='address'/state/countersname='name'/out-pkts/** is changed to /mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'**/state/counters name='name'/**. The

leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

- When the krt-nexthop-ack statement is configured, the RPD will wait for the next hop to get acknowledged by PFE before using it for a route. Currently, only BGP-labeled routes and RSVP routes support this statement. All other routes will ignore this statement.

Timing and Synchronization

- **Performance monitoring time interval with UTC on Junos OS Evolved platforms** —The performance monitoring (PM) time interval for 1-day bins on Junos OS Evolved platforms begins at midnight in the UTC zone, aligning with the standard behaviour of Junos OS. This synchronization allows you to maintain consistent performance monitoring schedules across platforms, enhancing the accuracy and reliability of network performance data.

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.
- **Load JSON configuration data with unordered list entries (ACX Series, PTX Series, and QFX Series)**— The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:

  - Use the `request system convert-json-configuration` operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.
  - Configure the `reorder-list-keys` statement at the `[edit system configuration input format json]` hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.

  When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

  [See json and request system convert-json-configuration.]

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.
- **Persistent CLI timestamps**—To have a persistent CLI timestamp for the user currently logged in, enable the `set cli timestamp` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements:

```
set system login class class name permissions permissions
set system login class class name cli timestamp set system login user username class class
name authentication plain-text-password
```

> **NOTE**: The default timestamp format is %b %d %T. You can modify the format per your requirements. For example, you can configure the following statement:
> `set system login class class name cli timestamp format "%T %b %d`

To enable timestamp for a particular user with default class permissions and format, configure the following statements:

```
set system login user username class class name authentication plain-text-password
set system login user username cli timestamp
```

Known Limitations

IN THIS SECTION

- General Routing | 20
- Routing Protocols | 22

Learn about known limitations in Junos OS Evolved Release 22.2R1 for ACX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- In some corner cases, traffic is not scheduled equally between strict priority queues. This can happen in the following scenario. Priority queue configured and is completely utilizing the bandwidth and remaining queues are starved and traffic completely dropping on those queues. In this state if we configure a second strict high priority queue traffic is not scheduled equally between strict priority queues. This is hardware specific issue, ACX7509 specific. If we have a shaper on priority queue this issue will not happen. Also if the traffic starts after the configurations no issues seen. PR1577035
- PTP to PTP noise transfer fails for frequencies 1. 0.03125 HZ 2. 0.123125 HZ. PR1608786
- On ACX7100-48L devices, syncE to PTP and syncE to 1pps noise transfer tests fails for frequencies 1. 0.00781 HZ 2. 0.01563 HZ 3. 0.03125 HZ 4. 0.06156 HZ 5. 0.12313 HZ. PR1608866
- The syncE to PTP and syncE to 1pps transient response marginally fails. This happens when the servo get the initial 100 nano seconds jump in one measurement window and the next 100 nano seconds in the next measurement window adjusting less initially. PR1608934
- On ACX7100-48L, enabling or disabling of PTP TC or BC causes all the interfaces to flap at the same time. PR1609927
- PTP to PTP noise transfer fails for frequency 0.03125 HZ. PR1611838
- On ACX7100-32C devices, the syncE to PTP and syncE to 1pps noise transfer tests fail for 1. 0.00781 HZ 2. 0.01563 HZ 3. 0.03125 HZ 4. 0.06156 HZ 5. 0.12313 HZ frequencies. PR1611911
- The `clear mpls lsp` operation is a destructive operation where it wipes off all existing routes and next-hops in the system and does a fresh reinstallation, the 10 seconds delay in traffic restoration for 16000 l3vpn routes might be attributed to programming delay in the hardware units combined with software model and the CPU capacity. PR1614413
- The learning rate of ACX7509 is same as ACX7100 when the host routes /128 routes are downloaded to LEM table in both ACX7100 or ACX7509. The PR is reporting an issue only when the scale goes beyond LEM table size. If the scale is within LEM table size, then the FIB downloading rate remains same in ACX7100 and ACX7509. PR1624365
- G.8275.1- G.8273.2 1PPS cTE performance test might be outside class-C when using channelized 10G ports for PTP BC on ACX7100-32C. On each reboot, the 1PPS cTE measurement might be within the class-C measurement threshold, or might randomly be out of it by a few nanoseconds. PR1629819

- G.8275.1- G.8273.2 1PPS cTE performance test might be outside class-C when using channelized 25G ports with 100G ports for PTP BC on ACX7100-32C. On each reboot, the 1PPS cTE measurement might be within the class-C measurement threshold, or might randomly be out of it by a few nanoseconds. PR1637268

- Ping and Traceroute works with reply mode as ip-udp (applicable to other Junos OS Evolved ACX series). Other reply mode (application-level-control-channel) works when we support BFD over VCCV. For ping, the default mode for MSPW Echo reply is application-level-control-channel. Hence, MSPW L2VPN ping needs reply-mode as ip-udp for the ping to work, as BFD over VCCV is not supported yet. For traceroute, the default mode is application-level-control-channel. Hence, MSPW L2VPN traceroute needs reply-mode as ip-udp for the traceroute to work, as BFD over VCCV is not supported. PR1642026

- Links might flap briefly (for a few milliseconds) if a switchover happens due to primary FEB power-fault. Workaround is to configure interface hold-timers at far-end routers. PR1652921

- 1. Why SIP and DIP increment case working fine for all cases (L3VPN, IPv4) other than 6VPE? Also all CCC and bridge family works fine when we increment SMAC and DMAC. For Ipv4 related services, there is a field called BRCM_HASH_FIELD_IPV4_ADDRESS which we are setting while configuring the BCM pipeline during hashing. Therefore, when there is a symmetrical increase in both Source IP and Destination IP, this flag takes care of load balancing. For IPV6 services, this field isnt supported as specified in the BCM documentation. Therefore on symmetrical increase in SIP and DIP, load balancing does not work. 2. Why 6VPE, it is not load-balanced? Reasoning from BCM. Its not that 6VPE service is not load-balanced. It is perfectly load-balanced as specified in the outputs below. xmen-s-p1b-d Seconds: 6 Time: 07:41:10 Interface Link Input packets (pps) Output packets (pps) ae0 Up 4798 (0) 4153317 (996) esi Up 0 (0) 0 (0) et-0/0/0 Up 0 (0) 0 (0) et-0/0/1 Up 2238 (0) 1983350 (797) <<== (Load balancing) et-0/0/2 Up 2560 (0) 2169967 (199) <<-(Load balancing) et-0/0/3 Up 0 (0) 0 (0) et-0/0/4 Up 0 (0) 0 (0) et-0/0/5 Up 4286778 (997) 8430 (8) Its just that for one particular case in which both are incrementing, the functionality doesnt work due to the reason provided in response 1. So we cannot say that the feature does not work. 3. In the 6PE case in SW set-up. The same behaviour of load balancing not working in 6PE case has been observed. PR1658411

- Junos OS Evolved follows Make-Before-Break (MBB) mechanism to program next-hop and route to achieve faster convergence. The mechanism installs new forwarding table entry before deleting old one, minimizing traffic loss during route convergence. However, it temporarily increases the number of forwarding paths programed in the Packet Forwarding Engine depending on number of times next-hop/route changes in short period of time. MBB is applied during link-flaps, graceful restarts (ldp), session flaps(ldp) etc. For deployments where the network device is running on the higher end of the tunnel scale limits, a link flap can easily exceed the scale of the device. Once Packet Forwarding Engine exceeds its forwarding table capacity any new nexthop add for a tunnel is ignored, resulting in traffic black hole for those NHs. A link flap though, will trigger MBB for only the tunnels associated with that particular link. If we take a worst-case situation that all the links flap at once and all the

tunnels are hence undergoing MBB, we have to keep the tunnel limit to half to be absolutely sure not to exceed the limit. PR1660472

Routing Protocols

- When NSR is enabled on routers with non-forwarding routing-instance having BGP peers, the BGP peers in that instance will not be successfully replicated to the backup RPD. PR1648707

Open Issues

IN THIS SECTION

- General Routing | 22
- EVPN | 25
- Infrastructure | 25

Learn about open issues in Junos OS Evolved Release 22.2R1 for ACX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- ACX7509: some of the interfaces from 16x100G and 20XSFP56 will not go down after evo-pfemand restart. PR1592388
- The error messages - **Supercon-core 0000:xx:xx:x: Supercon scratch**, are observed in the boot log of backup Routing Engine. The errors are observed only during booting of Backup Routing Engine and do not have any functional impact. These messages should be ignored.PR1594136
- Timing applications cannot be restarted successfully on the ACX7100 platform with the first software version G.8275.1.PR1597120

- With IGMPv3 reports received at a higher rate more than 1600 pps, packets are dropped due to control plane rate limit. Therefore, it is not possible to form 2,56,000 IGMP groups. Need to tweak DDOS configuration for reaching 2,56,000 IGMP groups. PR1599998
- The output of the `show system processes extensive` command shows high short term CPU utilization. The values ranges from 50 percent or higher for evo-pfemand. This is a single CPU view. As the ACX7509 system is a multi-core CPU, this has no impact on performance. PR1603899
- G.8275.1- G.8273.2 1PPS cTE performance test might be marginally outside class-C for PTP BC on ACX7100-48L. On each reboot, the 1PPS cTE measurement might be within the class-C measurement threshold, or might randomly be out of threshold by a few nanoseconds.PR1607381
- Interfaces mapped to the same Ethernet PHY flap when any one of the interfaces speed is modified. For example, the interface 0-7 flapping when speed(10g) on port 5 was removed or applied on the 20xsfp56 card. There can be a combination of speed as initial configuration. However, a link flap is observed within the ports of the port groups if the speeds of any port are reconfigured to other speeds. To avoid such a situation all the ports of the port group can have single-speed or do not reconfigure the speed within the port group.PR1608223
- The syncE to PTP and syncE to 1pps transient response marginally fails. This happens when the servo get the initial 100 nano seconds jump in one measurement window and the next 100 nano seconds in the next measurement window adjusting less initially. PR1611848
- On ACX7509 devices, 1GE interface does not comes up with copper 1G SFP-T optics and this issue is specific to copper 1G cables. PR1614286
- On ACX7509 devices, after multiple FPC online or offline, FPCs are going to fault state. PR1616227
- If a system is fully scaled across features and firewall is also scaled, CPU consumption might be more for a small window of around 5 seconds after every 18 seconds or so. Evo-pfemand might be busy collecting the scaled firewall statistics for that 5 second window and any other applications like 'pfe-cli' trying to execute commands might fail during it. PR1629342
- On ACX7100-32C and ACX7509 routers, 4x100g channels do not come up after multiple iterations Tx lase disabled alarm on. PR1631193
- On ACX7509 devices, on ungraceful removal of FEB, RCB crash/vmcore might be observed. RCB comes back up after reboot.PR1636780
- On ACX7509, on multiple FPC restart, link is not coming up with huge FEC errors. PR1639666
- With combination of triggers like restart rpd and fpc, route object leaks are observed in Packet Forwarding Engine. PR1641947
- When CLI based trigger of FPC restart is given and subsequently a Routing Engine switchover event is triggered, before the FPCs have come back online. The FPCs might not come online after switchover. It might be stuck in the offline or Fault state.PR1645305

- When the original flow egresses out through an aggregate Ethernet interface, the corresponding sampled sflow frame does not reflect the correct egress port number. This occurs only when the flow egresses out through an aggregate Ethernet interface, while it works fine in a non-aggregate Ethernet egress interface and the sflow frame reflects the correct egress port.PR1647870
- If FPC is restarted with multi-D scale configuration, then PICD or evo-pfemand app core will be observed occasionally. However, the app will be restarted and system will be restored automatically.PR1650302
- On ACX7100 and ACX7509 devices, OAM link fault management (LFM) discovery state is not correct for some interfaces and discovery state is either Active Send Local or Fault. PR1651580
- If the first port of a group used as channelised port (example port 8 from phy 3) instead of "unused" we don't see issue of particular phy going out of sync.PR1657531
- G.8275.1.enh 2Way cTE performance test might be outside class-C when using channelized 25G ports with 40G ports for PTP BC on ACX7100-32C. On each reboot, the 2Way cTE measurement might be within the class-C measurement threshold, or might randomly be out of it by a few nanoseconds.PR1662367
- In a dual routing engine setup, while switchover happens, there are some references that are help by the interfaces software which is not cleared. This anomaly is observed when switchover happens as a result of the primary routing engine/FEB being turned-off/rebooted. This memory leak does not have any impact on service and will be fixed in the subsequent release.PR1662411
- Interface flaps are observed on performing primary role switchover using the primary Routing Engine offline button press, however, no intermittent flaps are observed after switchover. PR1668509
- The traffic drop greater than 50 milliseconds is sometimes observed on an ungraceful jack out or in of primary Routing Engine, however, switchover happens and traffic resumes.PR1668780
- The traffic hit can sometimes be in seconds on a switchover caused by master FEB power-fault. The expected traffic hit is less than 50 milliseconds. There is no lingering effect due to this. The system will subsequently work normally.PR1668890
- The hwdre and evo-pfemand applications might crash if idmd, fabtoken and hwdre are restarted immediately after a FEB offline.PR1669130
- Switchover caused by primary RCB power-fault might cause links to go down.PR1669162
- In rare cases and with scaled scenarios, it is possible for Packet Forwarding Engine data path to not get properly programmed after a FEB offline or online. This can result in traffic not passing through. Workaround is to deactivate the scaled configuration, offline and online the FEB, and then activate the config again.PR1669211
- In a ACX7509 system with dual Routing Engine and FEB, when there is a power fault of primary FEB, a switchover should happen and backup Routing Engine and FEB will take up primary role. Post

switchover a VMCore might be generated in the new backup Routing Engine. This failure should not impact the system uptime. Post Vmcore backup Routing Engine will reboot and come back online. PR1671198

- If a FEB goes to fault state due to a power-fault (real or artificially triggered for testing), then the subsequent FEB offline can take a few minutes (instead of completing within a minute for a normal offline). There is no other collateral due to this. A FEB online subsequent to the delayed offline will work normally and the FEB will become fully functional again.PR1671719
- Links to backup FEB might go down on primary FEB restart. This can become a problem on a subsequent switchover. Workaround is to do an restart of the backup FEB after primary FEB has come back online (following the restart).PR1673274

EVPN

- On performing Graceful Routing Engine swithchover on ACX7509 device with EVPN services, some IFLs are missing. PR1646722

Infrastructure

- On triggerring Routing Engine switchover (graceful and ungraceful), large number of overlay and ref-counting ERROR messages are visible in the journalctl and syslog. These error messages start with prefix **OF: ERROR**. These error messages can be ignored. They do not impact the system behaviour. PR1639342

Resolved Issues

IN THIS SECTION

- Resolved Issues: 22.2R1 | **26**

Learn about the issues fixed in this release for ACX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

Resolved Issues: 22.2R1

IN THIS SECTION

- EVPN | 26
- General Routing | 26
- EVPN | 28
- Interfaces and Chassis | 28
- Platform and Infrastructure | 28
- Routing Protocols | 28
- User Interface and Configuration | 28

EVPN

- VxLAN encapsulation might fail to forward traffic to remote devices. PR1639204

General Routing

- Filter with forwarding-class and destination-class combined might not work. PR1595788
- On ACX7509 devices, the picd process generates core file when you restart FPC with 700,000 scaled MAC. PR1602352
- [ecmp] [ecmp] acx7509 : ::Picd core is observed when we perform restart FPC with 32000 ECMP groups scale. PR1609389
- On ACX7509 devices, multicast packets not shown in MAC statistics for LACP on LAG interfaces. Workaround, instead of querying MAC statistics under aggregate Ethernet interface, query for actual physical IFD interface created under aggregate Ethernet interface. PR1612105
- PTP performance might randomly become bad for 30 minutes. PR1614309
- On ACX7509 devices, Linux command to remove PCIe device causes connection to hang. PR1619368

- The issue is specific to slot-7, where slot-7, port-13 is not supported. As part of channelization constraints, both adjacent ports have to configure or both to non channelization. Since slot-7 port is unsupported, need to ignore channelization constraints for port the two ports.PR1620425
- On ACX7509, details of auto negotiation settings are not available on 1GigE interfaces PR1621991
- On ACX7509 devices, PCI device misses alarm raises for all FPCs from the backup Routing Engine. There are no functional impact. PR1627348
- After picd or rpdagent app restart multpile object-info anomalies for evo-pfemand are seen. PR1628843
- When a fan tray fails an alarm is raised. After rebooting the alarm gets cleared. However, the fan failure condition will be logged in the log file. PR1633353
- In a working and non-working logs, l2d index is different for VRRP group number 187. This is the same group for which packet is getting dropped out of 400 groups, other groups are working as expected. So there is some fix which went between working and NKWR related to l2dId which has exposed VRRP issue. Both VRRP MAC and interface MAC gets stored in SLU my_mac_hash table. For finding hash index for VRRP MAC we use l2dId. Protocol type and VRRP group number as a key. In a non-working scenario there is a collision between interface MAC and VRRP MAC on same hash index. Ideally hash movement should have happened to address collision but somehow it is not properly done. Going further we need to debug why hash movement is not happening and fix that code area. This code is very sensitive and requires a lot of testing before doing any changes. So we should commit it in DCB first before committing to any other release. PR1633986
- In scaled setup, while FPC restart, PICD is not releasing fruHwStatus status and causing PIC struck in online. PR1635941
- FEB jackout and jackin might generate a hwd core file while rebooting the system (but not at jackout or jackin time). There will not be any functionality issue as the core file is generated only during reboot sequence. System is expected to come up fine.PR1636243
- [interface] [interface] acx7509 : :: [46377: After Routing Engine primary switch with traffic during switchover, CPU utilization on new primary is more than 50% by EvoPfemand-main , java process and hence `show interfaces extensive` output is not displayed correctly for the first time after switchover] PR1636790
- [interface] [interface] acx7509 : :: [46377: in-octets, in-pkts, out-octets, out-pkts counters stats not replicated after switchover (without traffic flow, during switchover)] PR1641392
- On ACX7100 and ACX7509 devices, few control packets getting forwarded on ERPS discarding port also, leading to traffic loop. PR1641454
- [sw-evo-infra-utils] `show system core-dumps routing-engine both` CLI command does not display anything on ACX7509 devices. PR1646266

- TCAM field group entry shuffle is not traffic safe. PR1650266
- Rpd-agent core for high-scale of member nexthops in multicast composite nh. PR1640224
- There is a 15-20 seconds of iPv6 traffic loss after GRES. PR1655374

EVPN

- ACX7100-32C and ACX7100-48L, failed to forward some of the VXLAN encapsulated traffic to remote. PR1639204

Interfaces and Chassis

- On ACX7509 devices, there is a limitation in adding more than 64 member links in 1 LAG, whereas from ASIC, there is no limitation. PR1627951
- CFM sessions are not up after evo-pfemand restart or gets crashed. PR1634721

Platform and Infrastructure

- The line cards MPC10, MPC11, and LC-9600 might crash on Junos Evolved OS platform. PR1667716

Routing Protocols

- On ACX7509 routers, IS-IS, LPD, and BFD protocols flaps during graceful switchover while testing ldp oam. PR1638882

User Interface and Configuration

- The traffic might not flow after deleting or adding VLAN configuration with load override. PR1647853

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Devices

IN THIS SECTION

- What's New | 29
- What's Changed | 49
- Known Limitations | 53
- Open Issues | 54
- Resolved Issues | 58

These release notes accompany Junos OS Evolved Release 22.2R1 for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- What's New in 22.2R1 | 30

Learn about new features introduced in this release for PTX Series routers.

What's New in 22.2R1

IN THIS SECTION

- Chassis | 31
- Class of Service | 31
- EVPN | 32
- Interfaces | 32
- IP Tunneling | 33
- Juniper Extension Toolkit (JET) | 34
- Junos Telemetry Interface | 34
- MPLS | 37
- Multicast | 37
- Network Management and Monitoring | 38
- OpenConfig | 41
- Routing Protocols | 43
- Software Installation and Upgrade | 45
- Source Packet Routing in Networking (SPRING) or Segment Routing | 45
- VLANs | 46
- Additional Features | 46

To view features supported on the PTX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 22.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- PTX10001-36MR
- PTX10003
- PTX10004
- PTX10008
- PTX10016

The following sections highlight the key features in this release.

Chassis

- **Support for ASIC temperature-based voltage reduction (PTX10004, PTX10008, PTX10016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR)**—Starting in Junos OS Evolved Release 22.2R1, Junos OS Evolved dynamically manages the voltage of the DC-DC converters to keep ASIC temperature within safe operating limit.

  We've enabled voltage reduction by default as it also helps in reducing power consumption. To disable the feature, use the `set chassis fpc` *`slot`* `pfe` *`pfe-instance`* `temp-voltage-reduction-disable` command.

  [See fpc and show chassis fpc.]

- **Support for Fabric Hardening (PTX10004, PTX10008, and PTX100016 with PTX10K-LC1202-36MR line card)**—Starting in Junos OS Evolved Release 22.2R1, the PTX series routers support fabric hardening. Fabric hardening is a resiliency feature to detect fabric blackholing and attempt automatic recovery process to restore the Packet Forwarding Engines from blackhole condition. If recovery fails, interfaces are disabled to limit the blackholing and alarms are raised to indicate the unreachable PFEs.

  We've enabled fabric hardening by default.

  [See Fabric Hardening and Recovery and show chassis fabric reachability.]

Class of Service

- **Queue-depth monitoring support for virtual output queues (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Virtual output queue (VOQ) queue-depth monitoring, or latency monitoring, measures peak queue occupancy of a VOQ. Starting with Junos OS Evolved Release 22.2R1, PTX Series routers running Junos OS Evolved support VOQ queue-depth monitoring to report peak queue length for a given physical interface for each individual Packet Forwarding Engine.

  [See VOQ Queue-depth Monitoring.]

- **CoS support for EVPN-VXLAN (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 22.2R1, you can configure CoS features to prioritize traffic on an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) fabric.

  [See Implementing CoS on VXLAN Interfaces (Junos OS Evolved).]

- **Support for BGP flowspec DSCP action (PTX Series)**—Starting in Junos OS Evolved Release 22.2R1, you can configure mapping rules to translate a Differentiated Services code point (DSCP) value in the BGP flowspec (set DSCP action) into forwarding classes. With the mapping option enabled, you can change forwarding behavior, CoS treatment, and control EXP or DSCP rewrite of the packets. You can control rewrite of the packets using rewrite rules on a platform that does not support direct DSCP rewrite capability in a filter.

EVPN

- **EVPN-MPLS E-LAN flow-aware transport (FAT) label load balancing (PTX10001-36MR, PTX10004, PTX10008, PTX10016)** —Starting in Junos OS Evolved 22.2R1, you can configure provider edge (PE) devices to use FAT labels in an EVPN-MPLS routing instance, according to RFC 6391. Provider edge devices use these labels to load-balance EVPN-MPLS unicast packets across ECMP paths without needing to do deep packet inspection of the MPLS payload. This feature supports E-LAN with single-homing and multi-homing active/standby and active/active topologies and supports the VLAN-based, VLAN-bundle, and VLAN-aware bundle EVPN-MPLS variants.

  To enable load balancing using FAT labels in an `evpn` routing instance:

  - Configure the `flow-label-static` statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn` hierarchy level on PE devices to insert FAT flow labels into pseudowire packets sent to remote PE devices.

  - Configure the `flow-label` statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn` hierarchy level on PE devices to signal flow-label capability in the EVPN Layer 2 Attributes Extended Community by setting the flow-label (F) bit in the EVPN Type 3 route.

  [See flow-label and flow-label-static.]

- **Symmetric IRB with EVPN Type 2 routes (ACX7100, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can enable symmetric IRB EVPN Type 2 routing in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) edge-routed bridging (ERB) overlay fabric. With the symmetric routing model, leaf devices can route and bridge traffic on both ingress and egress sides of a VXLAN tunnel. Leaf devices use a transit VXLAN network identifier (VNI) and Layer 3 (L3) interfaces on the associated VLAN to exchange traffic across the VXLAN tunnels.

  We support this feature with `vlan-aware` and `vlan-based` MAC-VRF instance service type configurations. To enable this feature, you must also configure EVPN Type 5 routing with L3 VRF instances to establish intersubnet reachability among the EVPN devices.

  [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics and irb-symmetric-routing.]

Interfaces

- **Support for core-facing functionality (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we support core-facing functionality.

  [See core-facing.]

- **Support for media access control (MAC) accounting for source and destination macs for Layer 3 interfaces (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved

Release 22.2R1, we support media access control (MAC) accounting for source and destination macs for Layer 3 interfaces and aggregated Ethernet interfaces.

To enable MAC accounting, use the existing `mac-learn-enable` command under the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` or `[edit interfaces aex aggregated-ether-options ethernet-switch-profile]` hierarchy level.

[See show interfaces mac-database.]

- **Support for 400GbE port configuration commit check (PTX10003)**— Starting in Junos OS Evolved Release 22.2R1, we provide a validation check to restrict the configuration due to an invalid speed configuration. Earlier, you could configure any speed on any port and commit the configuration changes. Starting with this new validation check, whenever you commit an invalid configuration, the commit fails and throws an error message for the invalid configuration.

  [See Port Speed on PTX10003.]

- **Support for two 100GbE interface channelization on 400G-ZR-M (PTX10001-36MR)**—Starting in Junos OS Evolved Release 22.2R1, we support the channelization of two 100GbE interfaces on 400G-ZR-M.

  [See Port Speed on PTX10001-36MR Router Overview.]

IP Tunneling

- **Sharding support for dynamic IP-over-IP tunneling (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we support sharding for dynamic tunnels created as a result of BGP route resolution over a tunnel route. The protocol uses this tunnel route as a helper route for route resolution.

- **L2TPv3 header for load balancing (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, the device takes Layer 2 Tunneling Protocol version 3 (L2TPv3) header attributes into account when load balancing. When L2TPv3 packets are encapsulated in IP tunnels, the device identifies L2TPv3 packets by the IP protocol type 115. The session ID of the L2TPv3 packet headers distinguishes the traffic flows in L2TPv3-over-IP tunnels. The device uses the session ID of L2TPv3 packets in the hash computations to better manage traffic flows during load balancing.

  L2TPv3-based load balancing is enabled by default. To disable this feature, use the `no-tunnel-payload` statement at the `[edit enhanced-hash-key family any]` hierarchy.

  [See enhanced-hash-key.]

Juniper Extension Toolkit (JET)

- **JET API client gRPC connections limit (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, we have increased the maximum number of Juniper Extension Toolkit (JET) gRPC connections to 300. Use the `max-connections` statement at the `[edit system services extension-service request-response grpc]` hierarchy level to configure the maximum number of gRPC connections.

  [See max-connections.]

Junos Telemetry Interface

- **Chassis management error (cmerror) sensor support (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, use the resource path **/components/component/integrated-circuit/pipeline-counters/** to stream fine-grain, per-integrated circuit (IC) telemetry data that identifies the IC health and any packet drops or errors on the iC. This data is per router and per iC on a router. This feature supports the data model openconfig-platform-pipeline-counters.yang and is supported using Juniper proprietary gRPC remote procedure call (gRPC) and gRPC network management interface (gNMI).

  [See Telemetry Sensor Explorer.]

- **Forwarding table streaming support for non-default VRF (PTX10001-36MR, PTX10003, PTX10004, PTX10008 and PTX10016)**—Junos OS Evolved Release 22.2R1 extends support for forwarding information base (FIB, also known as forwarding table) streaming on JTI o include non-default virtual routing and forwarding (VRF) instances.

  [See Telemetry Sensor Explorer.]

- **Generic health and gNMI health sensors (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX 10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI streams telemetry health data that you can use to debug issues related to the telemetry subsystem.

  [See Telemetry Sensor Explorer.]

- **gRIBI sensor support (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI streams programmable routing protocol process (prpd) statistics and gRPC routing information base (RIB, also known as routing table) programming interface-related statistics.

  [See Telemetry Sensor Explorer.]

- **INITIAL_SYNC support for chassis management error (cmerror) telemetry statistics (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, Junos telemetry interface (JTI) supports INITIAL_SYNC for leaf nodes under subscription paths /junos/system/cmerror/configuration and /junos/system/cmerror/counters. This feature is supported using Juniper-proprietary gRPC and gNMI.

[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface.]

- **Integrated circuit sensors (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI provides periodic streaming and ON_CHANGE support for backplane capacity-related statistics for integrated circuits. This feature supports the data model openconfig-platform-integrated-circuit.yang. This feature is supported using Juniper-proprietary gRPC and gNMI.

  [See Telemetry Sensor Explorer.]

- **Network instance support enhancements (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.2R1, JTI supports new sensors for network instance statistics per OpenConfig modules openconfig-network-instance.yang and openconfig-routing-policy.yang. Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **OpenConfig interface configuration support (PTX10008 and PTX10016)**—Junos OS Evolved Release 22.2R1 enhances OpenConfig interface configuration support and state support.

  [See Mapping OpenConfig Interface Commands to Junos Configuration and Telemetry Sensor Explorer.]

- **OpenConfig interface configuration support for proxy ARP (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.2R1 introduces OpenConfig interface configuration support and state support for proxy ARP.

  [See Mapping OpenConfig Interface Commands to Junos Configuration and Telemetry Sensor Explorer.]

- **OpenConfig telemetry system model support (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130, and QFX5220)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig data models openconfig-telemetry.yang and openconfig-telemetry-types.yang. Support includes streaming of state data for dynamic subscriptions and configuration for persistent subscriptions.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **Optics data model support enhancements (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**—Junos OS Evolved Release 22.2R1 enhances support for optics based on the following OpenConfig data models:

  - openconfig-platform-transceiver.yang
  - openconfig-transport-types.yang
  - openconfig-terminal-device.yang

  Support includes OpenConfig configuration and streaming of operational state data.

[See Telemetry Sensor Explorer for state sensors and OpenConfig User Guide for configuration.]

- **Platform sensors (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, Junos Telemetry Interface provides periodic streaming and ON_CHANGE support (whenever applicable) for platform statistics. This feature supports the data model openconfig-platform.yang version 0.12.2. The feature is supported using Juniper proprietary gRPC and gNMI.

  [See Telemetry Sensor Explorer.]

- **Platform sensors for chassis components (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI supports operational state sensors for chassis backplace, fabric, fan, power supply, and storage components.

  [See Telemetry Sensor Explorer.]

- **Platform sensors for fan, line card, and power supply (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, JTI provides sensor streaming and ON_CHANGE support (whenever applicable) for platform state. This feature supports OpenConfig data models openconfig-platform-fan.yang version 0.1.1, openconfig-platform-linecard.yang version 0.1.2, and openconfig-platform-psu.yang 0.2.1. Juniper proprietary gRPC and gNMI support this feature..

  [See Telemetry Sensor Explorer.]

- **Two-Way Active Measurement Protocol sensor support (ACX7100-32C, ACX7100-48L PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, Junos telemetry interface (JTI) provides periodic streaming support for Two-Way Active Measurement Protocol (TWAMP) statistics. This feature supports IPv4 and IPv6 configurations for TWAMP managed sessions and TWAMP light sessions. This feature is supported using Juniper proprietary gRPC and gNMI.

  [See Telemetry Sensor Explorer.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

MPLS

- **Enhancements to BFD-triggered MPLS fast reroute (FRR) for unicast next hops and session-id-change-limiter-indirect (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700, and QFX5200)**—In Junos OS Evolved Release 22.2R1, we've enhanced the BFD-triggered fast reroute for unicast next hops and session-id-change-limiter-indirect to address the issue of traffic being silently discarded because of a session mismatch between the control plane and data plane.

  You can limit the re-programming of the number of parent nodes of the indirect-nexthop and prevent additional complexity in the Packet Forwarding Engine when the session-identifier id of the indirect nexthop is changed. To make those changes, use the `session-id-change-limiter-indirect` configuration statement at the [`edit routing-options`] hierarchy level.

  [See Bidirectional Forwarding Detection (BFD) for MPLS and session-id-change-limiter-indirect.]

- **Conditional RSVP label-switched path (LSP) metrics (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can configure conditional metrics for local statically configured LSPs. The conditional metrics are based on the dynamically changing interior gateway protocol (IGP) metric. Junos OS Evolved changes the LSP metric to the configured conditional metric that corresponds to the highest threshold the IGP metric reaches. You can configure up to four conditional metrics for an LSP.

  By default, the IGP metric of routes configured with the install statement is the IGP metric value of the LSP destination route. If you configure the `track-igp-metric <install-v4-prefixes> <install-v6-prefixes>` statement at the `[edit protocols mpls]` or `[edit protocols mpls label-switched-path lsp-name]` hierarchy level, the routes IGP installs use the IGP metric of the prefix instead.

  Use the `conditional igp-metric-threshold threshold-metric-value static-metric-condition-value` statement at the `[edit protocols mpls label-switched-path lsp-name metric]` hierarchy level to configure this feature. To check whether the conditional metric is configured, use the `show mpls lsp extensive` command.

  [See Configuring LSP Metrics, metric (Protocols MPLS), track-igp-metric (LSP), conditional-metric, and show mpls lsp extensive.]

Multicast

- **MSDP support for next-generation MVPN (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**— Starting in Junos OS Evolved Release 22.2R1, next-generation multicast virtual private network (MVPN) supports Multicast Source Discovery Protocol (MSDP) in virtual routing and forwarding (VRF) instances.

  [See Understanding Next-Generation MVPN Control Plane.]

Network Management and Monitoring

- **Support for the gRPC Network Operations Interface (gNOI) `CertificateManagement` service (PTX10008 and PTX10016)**—Starting in Release 22.2R1, Junos OS Evolved devices support gRPC Network Operations Interface services. You can execute supported CertificateManagement service remote procedure calls (RPCs) to manage certificates on the network device. Using gNOI operations enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

  [See gNOI Certificate Management Service.]

- **Support for the gRPC Network Operations Interface (gNOI) `diag` service (PTX1004, PTX10008, and PTX10016)**—Starting in Release 22.2R1, Junos OS Evolved devices support gRPC Network Operations Interface services. You can execute supported Diag service remote procedure calls (RPCs) to perform diagnostic operations such as bit error rate tests (BERTs) on the network device. Using gNOI operations enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

  [See gNOI Diagnostic (Diag) Service.]

- **Support for the gRPC Network Operations Interface (gNOI) `file`, `os`, and `system` services (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Release 22.2R1, Junos OS Evolved devices support gRPC Network Operations Interface services. You can execute supported File, OS, and System service remote procedure calls (RPCs) to perform common file and system operations on the network device. Using gNOI operations enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

  [See gNOI File Service, gNOI Operating System (OS) Service, and gNOI System Service.]

- **Support for YANG metadata annotations for configuration operations (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we provide the `junos-configuration-metadata` YANG module. This module defines metadata annotations that you can use to perform specific configuration operations in YANG-compliant NETCONF sessions. Supported operations include adding comments to the configuration, deactivating and activating configuration statements, and protecting and unprotecting configuration statements.

  [See YANG Metadata Annotations for Junos Devices .]

- **Support for additional family in port mirroring (PTX10001-36MR; LC1201 and LC1202 line cards in PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can configure family `any` (as well as the earlier family options, `inet` and `inet6`) for local port mirroring and remote port mirroring:

  - Family `any` (for family `any`, `ccc`, `ethernet-switching`, or `mpls`)

NOTE: You use the family `any` configuration option to process all 4 families.

As of Release 22.2R1, you no longer configure port mirroring by using the `[edit forwarding-options port-mirroring analyzer]` hierarchy on the PTX devices.You now use `[edit forwarding-options port-mirroring]` for local port mirroring or `[edit forwarding-options port-mirroring instance instance-name]` for remote port mirroring, with both of those configurations also requiring a firewall filter.

The following configuration statements are no longer part of the port mirroring configuration on PTX: `next-hop` for family `any`; family `vpls`; `no-filter-check`; `hosted-service`; `server-profile`.

[See Example: Configure Port Mirroring with Family any and a Firewall Filter.]

- **gRIBI `Flush()` remote procedure call (RPC) support (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, these devices support the gRPC Routing Information Base Interface (gRIBI) `Flush()` RPC. The `Flush()` RPC removes all the server's gRIBI programmed routes that match what is described in the `FlushRequest` message. Sending a `Flush()` request is a quick and easy way to delete gRIBI programmed routes from the server.

  [See gRIBI.]

- **gRIBI API enhancements (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we offer enhanced support for the gRPC Routing Information Base Interface (gRIBI) API including:

  - The gRIBI client only programs an `IPv4Entry` AFT entry on the server after it receives acknowledgments from the server that the server received the associated `NextHopGroup` and `NextHop` messages. If the client programs an `IPv4Entry` AFT entry on the server without the `NextHopGroup` message, it adds the route to the server as a hidden route.

  - You can now perform hierarchical lookups by using the `IPv4Entry` AFT entry to program IP-IP tunnel endpoints and site group virtual IP address routes.

  - We now support arbitration when multiple clients are connected to the gRIBI server.

  [See gRIBI.]

- **gRIBI RIB-FIB acknowledgment (PTX10008)**—Starting in Junos OS Evolved Release 22.2R1, the Packet Forwarding Engine sends an acknowledgment when you successfully program a route in the Packet Forwarding Engine using the gRPC Routing Information Base Interface (gRIBI) API. When the gRIBI API fails to program a route in the Packet Forwarding Engine after a timeout period passes, the Packet Forwarding Engine sends an error message. You can configure the length of this timeout. The acknowledgment is only valid for the most recent route. If an older route sends an acknowledgment but the new route does not, the Packet Forwarding Engine records the acknowledgment as an error.

  Use the `show route extensive` command to display the acknowledgment status. The acknowledgment status is persistent across rpd process restarts.

[See gRIBI.]

- **gRIBI API route reconciliation and recovery after client reconnection (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we have enhanced the gRPC Routing Information Base Interface (gRIBI) API to better recover routes after the gRIBI server or rpd process goes down. When the client reconnects to the server, the client automatically sends a gRIBI `Get()` RPC request to the server. If GRES is configured, the client reconciles the routes on the server. If the client sends another `Get()` RPC request, the answering `GetResponse` response stream includes the active reconciled routes on the server.

  If you configure GRES but do not configure non-stop routing, the gRIBI API also recovers routes after a Routing Engine switchover.

  [See gRIBI.]

- **gRIBI support for fallback to default route in VRF instance (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can use the gRPC Routing Information Base Interface (gRIBI) API to program a default route in a traffic engineering virtual routing and forwarding (VRF) instance as the backup route. Traffic falls back to this backup route when all other routes are unavailable, which prevents traffic loss. This default route has a next hop with decapsulation and looks up routes in the default VRF. Add the default route to the VRF first so the future routes you configure in the VRF will use it as a fallback route.

  [See gRIBI.]

- **gRIBI enhancements for IP-in-IP tunneling (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, the gRPC Routing Information Base Interface (gRIBI) API supports IP-in-IP encapsulation. Set the `backup_next_hop_group` field in the `NextHopGroup` message to perform decapsulation and look up routes in other routing instances. You can program a service route using multiple IP-IP tunnel endpoints.

  [See gRIBI.]

- **gRIBI support for automatic tunnel fallback (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can use the gRPC Routing Information Base Interface (gRIBI) API to program backup tunnels for automatic fallback. Use the gRIBI `Modify()` remote procedure call (RPC) to program a backup next-hop group. If the primary tunnel goes down, the device automatically switches the traffic to this backup tunnel. Automatic fallback quickly reroutes traffic to reduce traffic loss. This feature supports IPv4 transport for dynamic IP-IP tunnels with an IPv4 or IPv6 payload.

  Use the `PolicyForwardingEntry` message to program policy-based forwarding on the gRIBI server. Policy-based forwarding ensures that traffic moved to the backup tunnel remains in the tunnel regardless of what the routing table says.

  [See gRIBI.]

OpenConfig

- **Network instance support enhancements (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**—Junos OS Evolved Release 22.2R1 supports new sensors for network instance statistics based on OpenConfig (OC) modules **openconfig-network-instance.yang** and **openconfig-routing-policy.yang**. Support includes OC configuration and streaming of operational state data.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **OpenConfig authentication, authorization, and accounting (AAA) configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX 10016)**—Junos OS Evolved Release 22.2R1 introduces AAA configuration support for remote procedure calls (gRPC) RPCs and the gRPC-based services gRPC Network Operations Interface (gNOI), gRPC Routing Information Base Interface (gRIBI), and gRPC Network Management Interface (gNMI).

  [See Mapping OpenConfig AAA Commands to Junos Operation.]

- **OpenConfig firewall filter configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.2R1 introduces firewall filter configuration support for gRPC.

  [See Mapping OpenConfig Firewall Filter Commands to Junos Configuration.]

- **OpenConfig interface configuration support for proxy ARP (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.2R1 introduces OpenConfig interface configuration support and state support for proxy ARP.

  [See Mapping OpenConfig Interface Commands to Junos Configuration and Telemetry Sensor Explorer.]

- **OpenConfig interface configuration support (PTX10008 and PTX10016)**—Junos OS Evolved Release 22.2R1 enhances OpenConfig interface configuration support and state support.

  [See Mapping OpenConfig Interface Commands to Junos Configuration and Telemetry Sensor Explorer.]

- **OpenConfig IS-IS configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.2R1 introduces IS-IS configuration support.

  [See Mapping OpenConfig ISIS Commands to Junos Configuration .]

- **OpenConfig MPLS label distribution protocol (LDP) configuration support (ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)—**—Junos OS Evolved Release 22.1R1 introduces OpenConfig (OC) support for LDP configuration based on the OpenConfig data model openconfig-mpls-ldp.yang.

  [See Mapping OpenConfig MPLS Commands to Junos Configuration.]

- **OpenConfig Proxy ARP and IPv6 duplicate address detection (DAD) configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.1R1 introduces OpenConfig support for proxy ARP and IPv6 DAD configurations based on the OpenConfig data model openconfig-if-ip.yang version 3.0.0.

  [See Mapping OpenConfig Interface Commands to Junos Configuration .]

- **OpenConfig telemetry system model support (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130, and QFX5220)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig data models openconfig-telemetry.yang and openconfig-telemetry-types.yang. Support includes streaming of state data for dynamic subscriptions and configuration for persistent subscriptions.

  [See Telemetry Sensor Explorer for telemetry support and Mapping OpenConfig Telemetry System Model Commands to Junos Configuration for configuration.]

- **Optics data model support enhancements (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 22.2R1 enhances support for optics based on the following OpenConfig data models:

  - openconfig-platform-transceiver.yang
  - openconfig-transport-types.yang
  - openconfig-terminal-device.yang

  Support includes OpenConfig configuration and streaming of operational state data.

  [See Telemetry Sensor Explorer for state sensors and OpenConfig User Guide for configuration.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

- **P4 runtime configuration for Layer 2 support (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can configure a P4 runtime agent for Layer 2 support. The P4 runtime agent runs as a Junos OS Evolved application called p4-switch on the Routing Engine and implements a P4 runtime service over a gRPC channel. The P4 runtime agent registers with the

Juniper Extension Toolkit (JET) services daemon (jsd) and opens the gRPC connections by using JSD to listen for P4 requests from clients.

You configure the JSD port with the command `set services extension-service request-response grpc clear-text port port-number max-connections number`. You use OpenConfig to configure that port as a p4rt port. You then create a Layer 2 and Layer 3 interfaces configuration on the device with VLAN tags and the IP address for the device, allowing Layer 2 packets to be punted and injected between the device and the SDN controller. You use firewall filters to direct packet punts to the p4-switch.

[See OpenConfig User Guide.]

Routing Protocols

- **Support for IS-IS flood-reflector interfaces (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we support the IS-IS flood reflector that offers better scalability for a Level 2 topology. Flood reflectors enable the creation of topologies where Level 1 areas provide transit forwarding for Level 2 destinations within a Level 2 topology.

  You can configure a node to be either the reflector or the client. To enable the reflector, include the `flood-reflector reflector cluster-id` statement at the `[edit protocols isis level level number]` hierarchy level.

  To enable the flood reflector client, include the `flood-reflector client` statement at the `[edit protocols isis level level number]` hierarchy level.

  The flexible tunnel interfaces (FTIs) are designated as flood-reflector interfaces.

  - On a flood-reflector node:

    To enable the flood reflector on an FTI, include the `flood-reflector` statement at the `[edit protocols isis interface interface name level level number` hierarchy level.

  - On a flood-reflector client node:

    To enable the flood reflector on an FTI, include the `flood-reflector cluster-id` statement at the `[edit protocols isis interface interface name level level number]` hierarchy level.

  NOTE: You can configure the flood reflector on FTIs at Level 2 only.

  [See How to Configure Flood-Reflector Interfaces in IS-IS Networks.]

- **Migration to IS-IS flood reflector interfaces and RSVP Explicit Route Object (ERO) expansion (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we support the migration of IS-IS flood reflector and RSVP-Traffic Engineering (TE) ERO in a

CLOS fabric network. This feature allows the creation of IS-IS flood reflectors and RSVP TE label-switched paths (LSPs) with minimal traffic loss.

[See How to Configure Flood-Reflector Interfaces in IS-IS Networks.]

- **Nonstop active routing (NSR) support with BGP RIP sharding and BGP UpdateIO features (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we've enabled nonstop routing (NSR) for BGP RIP sharding and BGP UpdateIO features. With NSR enabled, the backup Routing Engine and backup routing protocol process (rpd) become the primary Routing Engine without negatively affecting the BGP peering sessions with the neighbors if the primary Routing Engine fails. The backup rpd processes the replicated BGP control-plane information and populates the route state in the same multithreaded manner as in the primary rpd.

  After you configure NSR, the `show bgp neighbor` and `show bgp summary` commands display the information about the specific shards in the backup Routing Engine. To display the replicated information for a specific shard in the `show bgp replication` command, use the `rib-sharding` *`shard-name`* option.

  See [show bgp neighbor, show bgp summary, show bgp replication, and BGP Overview.]

- **TCP authentication option (TCP-AO) for resource public key infrastructure (RPKI) validation sessions (ACX7100-32C, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use the TCP authentication option to authenticate RPKI validation sessions for securing the Internet's routing infrastructure, such as the BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

  To enable a TCP authentication option chain to authenticate an RPKI validation session, use the configured `authentication-algorithm` **`ao`** `and authentication-key-chain` *`keychain`* at the [`edit routing-options validation group` *`group_name`* `session` *`address`*] and [`edit routing-options validation group` *`group_name`*] hierarchy level.

  [See TCP Authentication Option (TCP-AO)]

- **BGP extended route retention (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, we have enhanced the long-lived graceful restart (LLGR) capabilities for BGP Helper and introduced the new configuration statement `extended-route-retention` at `[edit protocols bgp group neighbor graceful-restart long-lived]` hierarchy level. When you enable this feature, `extended-route-retention` supports LLGR helper mode regardless of the BGP peer LLGR capabilities. We have also updated the output of the following existing operational commands:

  - show bgp neighbor
  - show route extensive

  [See graceful-restart-long-lived-edit-protocols-bgp.]

- **RPD Object Reference count Anomaly checker (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**— Starting in Junos OS Evolved Release 22.2R1, we introduce a generic reference count infrastructure that all the modules in the routing protocol process (rpd) can use. The module maintains lock and unlock statistics corresponding to each object type in use. Any application can call the refcount increment or decrement API when an object is referred or deferenced. The module also provides a mechanism to detect anomalies such as a leak or overflow in an object's refcount.

Software Installation and Upgrade

- **Support for DHCPv6 on zero-touch provisioning (ZTP) (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, ZTP supports the DHCPv6 client on management and WAN interfaces. During the bootstrap process, the device first uses the DHCPv4 client to request information regarding image and configuration files from the DHCP server. The device checks the DHCPv4 bindings sequentially. If one of the DHCPv4 bindings fails, the device continues to check for bindings until provisioning is successful. If no DHCPv4 bindings exist, however, the device checks for DHCPv6 bindings and follows the same process as for DHCPv4 until the device can be provisioned successfully. The DHCP server uses DHCPv6 options 59 and 17 and applicable sub-options to exchange ZTP-related information between itself and the DHCP client.

  See [Zero Touch Provisioning].

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for application-specific link attribute in OSPFv2 for segment routing traffic engineering (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, you can advertise different te-attributes such as te-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link. This is done using flexible algorithm specific application-specific link attribute as defined in RFC 8920.

  To configure flexible algorithm application-specific te-attribute, include the `application-specific` statement at the `[edit protocols ospf area interface]` hierarchy level and the `strict-asla-based-flex-algorithm` statement at the `[edit protocols ospf source-packet-routing]` hierarchy level.

  [See Understanding OSPF Flexible Algorithm for Segment Routing.]

- **BGP Classful Transport (CT) support for IPv6 and Segment Routing Traffic-Engineered (SR-TE) color-only support (PTX10001-36MR, PTX10003, and PTX10008)**— Starting in Junos OS Evolved Release 22.2R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have also enhanced the `transport-class` configuration statement to provide strict resolution without falling back on best-effort tunnels.

[See use-transport-class, BGP Classful Transport (BGP-CT) with Underlying Colored SR-TE Tunnels Overview.]

VLANs

- **Support for no-local-switching functionality (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 22.2R1, we support no-local-switching functionality .

  [See no-local-switching.]

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **Accept BGP routes with accept-own community** (ACX7100-32C, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See BGP accept-own Community and accept-own.]

- **Adaptive RSVP update-threshold for optimal IGP flooding** (PTX10001-36MR, PTX10002, PTX10003, PTX10004, PTX10008, and PTX10016)—You can now configure the RSVP update threshold percentage and threshold value to adaptively pace IGP updates.

  [See update-threshold.]

- **Collect ON_CHANGE BGP RIB telemetry statistics** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5220)

  [See Telemetry Sensor Explorer.]

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding** (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)

  [See Telemetry Sensor Explorer.]

- **Inline active flow monitoring using IP Flow Information Export (IPFIX) and version 9 templates** (PTX10003)—We now support:

  - IPv4 and IPv6 traffic on IRB interfaces.
  - A BGP next-hop address in the IPv6 and MPLS-IPv6 templates: Information Element 63, IPv6 BGP NextHop Address, is now available.

  [See Inline Active Flow Monitoring on IRB Interfaces and IPFIX and Version 9 Templates.]

- **Inline monitoring services (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers with the JNP10K-LC1201 or JNP10K-LC1203 linecards)**—

  Inline monitoring services enable you to monitor different streams of traffic at different sampling rates on the same interface. The inline monitoring services in the Junos OS Evolved implementation show the following differences from the inline monitoring services in the Junos OS implementation:

  - A maximum of seven inline-monitoring instances
  - A maximum clip length of 256 bytes starting from the Ethernet header
  - Sampling rate per instance only
  - Port-mirroring and inline-monitoring firewall filter actions on different terms
  - Some inline-monitoring-services configuration options are available only in the Junos OS implementation, such as the options to specify the DSCP bits, forwarding class, routing instance, template ID, and option template ID.

  [See Inline Monitoring Services Configuration.]

- **IP-IP tunnel stitching** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation and Example: Configuring Next-Hop-Based IP-Over-IP Dynamic Tunnels.]

- **IPFIX IPv4 and IPv6 template support for forwarding class and packet loss priority (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**

  [See IPFIX and Version 9 Templates.]

- **Juniper Resiliency Interface (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers with the JNP10K-LC1201 or JNP10K-LC1203 linecards)**

  [See Juniper Resiliency Interface.]

- **Next-Generation MVPN** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)—We now support:

  - Rendezvous-point tree (RPT)-shortest-path tree (SPT) mode.
  - Restart individual PFE instances.
  - Turnaround provider edge (PE) device.
  - RP Mechanisms, including auto rendezvous point (RP), bootstrap router (BSR), and embedded RP.

  [See Understanding Next-Generation MVPN Control Plane.]

- **NSR Support for RIP and RIPng** (PTX10004, PTX10008, and PTX10016)

[See Nonstop Active Routing Concepts.]

- **Per Packet Forwarding Engine-Restart support (PTX10004, PTX10008, and PTX10016)**—

  Support includes:

  - Disable, offline, and restart of individual Packet Forwarding Engines.
  - Filter-based GRE encapsulation.
  - Segment routing–traffic engineering (SR-TE) per path telemetry statistics support, including ingress IP/binding-segment identifiers (SIDs) and MPLS transit traffic.
  - Flow Aware Transport (FAT) label transport for Layer 2 VPN.
  - VLAN tag manipulation support for CCC.
  - Segment routing over UDP.
  - Multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows.

  [See fpc (M320, T320, T640, and PTX Series Routers), Understanding Filter-Based Tunneling Across IPv4 Networks, source-packet-routing, flow-label-transmit, Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough, Next-Hop-Based Dynamic Tunnels, and Understanding Multicast-Only Fast Reroute.]

- **Support for BGP Classful Transport (CT) with underlying colored Segment Routing Traffic-Engineered (SRTE) tunnels (PTX10001-36MR, PTX10003, PTX10004, PTX10008 and PTX10016)**—BGP-CT can resolve service routes using the transport routing information bases (RIB)s, also known as routing tables and compute the next hop. Services currently supported over BGP-CT can also use the underlying SRTE colored tunnels for route resolution.

  [See use-transport-class.]

- **Support for EVPN-virtual private wire service (VPWS) (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—We've extended support for EVPN-VPWS to the listed PTX Series routers. By default, control-word is enabled on these platforms. To disable the control word feature, use the `set routing-instances` *`routing-instance-name`* `protocols evpn no-control-word` command.

  [See Overview of VPWS with EVPN Signaling Mechanisms.]

- **Support for firewall filtering using flood policer, IRB, and service provider egress filtering** (PTX10003)

- **Support for match conditions next-header and payload-protocol to filter IPv6 packets based on Next Header (NH) values of first and last headers** (PTX10003)

  The next-header matches on intermediate extension headers are not supported. The payload-protocol match is applicable only if the first extension header is Hop-by-Hop.

[See Firewall Filter Match Conditions for IPv6 Traffic]

- **Supported transceivers, optical interfaces, and direct attach copper (DAC) cables** —Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update this tool and provide the first supported release information when the optic becomes available.
- **Translation of multicast virtual private network (MVPN) Type 5 routes to Multicast Source Discovery Protocol (MSDP) SA routes** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See Converting MVPN Type 5 Routes to MSDP SA.]

What's Changed

IN THIS SECTION

- What's Changed in Release 22.2R1-S1 | 49
- What's Changed in Release 22.2R1 | 50

Learn about what changed in these releases for PTX Series routers.

What's Changed in Release 22.2R1-S1

IN THIS SECTION

- Juniper Extension Toolkit (JET) | 49
- Junos OS API and Scripting | 50

Juniper Extension Toolkit (JET)

- **Python 3 is the default and only Python version for executing Juniper Extension Toolkit Python scripts (ACX Series, PTX Series, and QFX Series)**—Junos OS Evolved supports only Python 3 for executing Juniper Extension Toolkit (JET) scripts written in Python. Python 2.7 is no longer supported for executing JET scripts, and we've deprecated the `language python` statement at the `[edit system scripts]` hierarchy level.

[See Understanding Python Automation Scripts for Junos Devices.]

Junos OS API and Scripting

- **Deprecated functions in the `libpyvrf` Python module (ACX Series, PTX Series, and QFX Series)**—The `libpyvrf`Python module no longer supports the `get_task_vrf()` and `set_task_vrf()` functions.

  [See How to Specify the Routing Instance in Python 3 Applications on Devices Running Junos OS Evolved.]

What's Changed in Release 22.2R1

IN THIS SECTION

- General Routing | 50
- High Availability | 50
- Interfaces and Chassis | 51
- OpenConfig | 51
- Network Management and Monitoring | 51
- Routing Protocols | 51
- User Interface and Configuration | 52
- VPNs | 52

General Routing

- Change in in unnumbered-address support for GRE tunnel?Starting in Junos OS Release 24.4R1, there is a behavioural change in unnumbered-address support for GRE tunnel with IPV6 family and display donor interface for both IPV4 and IPV6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level.

  See [show interfaces.]

High Availability

- **"Switchover Status Ready" incorrectly describes the status of the backup Routing Engine (RE) after node reboot (PTX10004, PTX10008, PTX10016)**— During preparation for switchover between master RE and backup RE running Junos OS Evolved releases prior to 22.2R1, "Switchover Status

Ready" from the `show system switchover` command on the backup RE node, after system reboot, incorrectly describes the status of the backup RE.

The incorrect status description results from a discrepancy between the master RE and the backup RE both using local uptime to determine if sufficient time had elapsed before declaring "Switchover Status Ready".

Use the `request chassis routing-engine master switch` command on the master RE and the backup RE to obtain the correct status when preparing for switchover.

[See show system switchover and request chassis routing-engine master.]

Interfaces and Chassis

- **JNP10K-PWR-DC2 power supplies installed in PTX10008 and PTX10016 routers display as online when the power supplies are switched off**— JNP10K-PWR-DC2 power supplies installed in PTX10008 and PTX10016 routers in which Junos OS Release 21.4R1 or Junos OS Evolved Release 21.4R1 is installed display as online in the output of the command `show chassis environment psm` when the input power feeds are connected, but the power switch on the power supplies are switched off.

- **Performance monitoring time interval with UTC on EVO platforms**— The performance monitoring (PM) time interval for 1-day bins on EVO platforms begins at midnight in the UTC zone, aligning with the standard behaviour of Junos. This synchronization allows you to maintain consistent performance monitoring schedules across platforms, enhancing the accuracy and reliability of network performance data.

OpenConfig

- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".

Network Management and Monitoring

- **Changes to the NETCONF `<edit-config>` RPC response (ACX Series, PTX Series, and QFX Series)**— When the `<edit-config>` operation returns an error, the NETCONF server does not emit a `<load-error-count>` element in the RPC response. In earlier releases, the `<edit-config>` RPC response includes the `<load-error-count>` element when the operation fails.

Routing Protocols

- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the `[edit system services ssh]` hierarchy level.

In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the `[edit system services ssh]` hierarchy level.

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.
- **Persistent CLI timestamps**—To have a persistent CLI timestamp for the user currently logged in, enable the `set cli timestamp` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements:

```
set system login class class name permissions permissions
set system login class class name cli timestamp set system login user username class class
name authentication plain-text-password
```

NOTE: The default timestamp format is %b %d %T. You can modify the format per your requirements. For example, you can configure the following statement:
`set system login class class name cli timestamp format "%T %b %d`

To enable timestamp for a particular user with default class permissions and format, configure the following statements:

```
set system login user username class class name authentication plain-text-password
set system login user username cli timestamp
```

VPNs

- **Changes to `show mvpn c-multicast` and `show mvpn instance` outputs** — The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

  See [show mvpn c-multicast.]

- **Enhancement to snmp mib command behavior (PTX10008)**—Starting in Junos OS Evolved, when you execute show snmp mib walk decimal command, the output parameter

jnxRedundancySwitchoverReason is not working as expected, which always show the value 0 instead of expected values. Now, jnxRedundancySwitchoverReason output parameter is corrected to expected behavior with the following expected values. `jnxRedundancySwitchoverReason OBJECT-TYPE SYNTAX INTEGER`

See [show snmp mib.]

- **Support for Embedded RP on PTX10008**—From this release, we support the Embedded RP feature on PTX10008 devices.

  See [Configuring Embedded RP.]

Known Limitations

IN THIS SECTION

- General Routing | 53
- User Interface and Configuration | 54

Learn about known limitations in Junos OS Evolved Release 22.2R1 for PTX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- On Junos OS Evolved PTX10008 platforms, if multiple SIBs are in offline state and if GRES is performed simultaneously, the SIBs might be stuck in offline state for sometime. PR1554423
- If multiple SIBs are in offline state and halt the primary Routing Engine, the SIBs might be stuck in the offline state for sometime, before it goes to offline state. PR1584712
- When an interface level shapers are applied, it might not draw accurate scheduler accuracy for all the queues under a port. PR1657883

User Interface and Configuration

- In a dual Routing Engine system, while ZTP is in progress and configuration commit happens on the primary Routing Engine RE0 (as part of ZTP) and configuration commit is not yet completed on RE0, if Routing Engine switchover is triggered then configuration commit is aborted or stopped ungracefully on the node RE0. The ZTP process starts freshly on the new primary node RE1. Due to the ungraceful exit of commit process on RE0, the ZTP might not succeed on RE1 because it can face commit failures on the primary node RE1. PR1649786

Open Issues

IN THIS SECTION

- General Routing | **54**
- Interfaces and Chassis | **57**
- Infrastructure | **57**
- Network Management and Monitoring | **57**
- Routing Policy and Firewall Filters | **58**
- User Interface and Configuration | **58**

Learn about open issues in Junos OS Evolved Release 22.2R1 for PTX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- The fabsopke-fchip core file can be seen if fabsopke-fchip restarts and sibs go offline simultaneously. Any previous alarm does not get cleared. PR1525577
- Sflow ingress sampling reports incorrect OIF and NH with MPLS traffic (MPLS-IPV4/MPLS-IPV6 with ECMP and Ultimate Hop Pop(UHP)) on the egress node. PR1602448

- 10,000 Term ISF filter (with or without fast-lookup filter configuration statement) with a set of add or delete ISF events leads to an evo-aftmand core file. This is not an FFT crash as the filter is programed in FLT (hence a baseline issue). PR1610506

- The linux kernel in Junos OS Evolved puts an autoflowlabel on every IPv6 packet. This flow label is transparent to daemon process, which uses a null value for it and calculates the NH details. Packet Forwarding Engine however takes the flow label into account and calculates the NH details. This difference in calculation of NH details leads to a mismatch in the path the packet takes to the destination and can cause traceroute to fail. PR1618406

- In some reload scenarios, an alarm **Major Application config-sync fail on node Re1** indicates **config-sync** failure, that can cause configuration mismatch. Work around : Restart the service from CLI using the command `restart config-sync` when the alarm is cleared. PR1629952

- PTX10001-36MR: 400G ZR-M 740-131169 modules do not come up in odd ports due to crossing temperature threshold which results into optics shutdown. Manually configuring fan speeds to 100% does help in bringing the optics up. PR1631279

- The workaround is to deactivate or activate auto-rp configuration on PTX10003-80C and PTX10003-160C. PR1634982

- System reboot or boot up with traffic could result in init time fabric link crc errors and cause traffic drop.PR1635178

- The issue is 400G-ZR/400G-ZR-M optic transceiver firmware upgrade failure. It is a generic issue faced across all single RU platforms supporting QDD-400G-ZR/QDD-400G-ZR-M. PR1638284

- Physical LED status is functioning as expected; However, the MIB state reflection of the physical LEDs is not as expected. PR1642816

- All the leafs under backplane-facing-capacity is exported per Packet Forwarding Engine or Fabric Packet Forwarding Engine ASIC. A component is identified with key /components/component[name] for **consumed-capacity** leaf, component name can be either PfeE name or FabricPfeE name. For all other leafs, key is always with FabricPfeE name. PfeE name example: **/Chassis[0]/Fpc[0]/Pfe[0] FabricPfeE name example: /Chassis[0]/Fpc[0]/Fchip[0]**]. All other leafs are exported with FabricPfeE name as key. Key Example: **(only for consumed-capacity) kv { key:__prefix__, str_value:/components/component[name='/Chassis[0]/Fpc[0]/Pfe[0]']/ } On release 22.2R1-EVO onwards "consumed-capacity" and all other leafs will be exported with FabricPfeE name as key. Key Example: (For all leafs) kv { key:__prefix__, str_value:/components/component[name='/Chassis[0]/Fpc[0]/Fchip[0]']/** PR1648354

- Next-header match in IPv6 firewall filter does not work as expected. Next-header matches the payload-protocol (last-header) on Junos OS Evolved PTX Series platform. PR1648864

- Filter is not hit, after trying to program unsupported combination of match types in filter configuration.PR1648923

- Junos OS Evolved: PTX10003 : Unsupported bit-op-type message seen for tcp flag match - **"(syn & ack) !(syn & ack & rst)"**. PR1649253
- Sensors under /components/component/integrated-circuit/backplane-facing-capacity/state/ are not exported in UDP through WAN port but exporting through management port works. PR1649876
- PTX10008: Link fails to come up with broken DOM after channelization. PR1654590
- When the 2x100G channelization is configured on the interface, and we have a ZR-M plugged into that port, when a hotswap is done with a 2x100G LR4, the LR4 does not come up. The port can be set to **unused** and then the unused configuration can be deleted to bring the LR4 links up. PR1655180
- When disabling or removing warm standby configuration and enabling or adding NSR, split these into two separate configuration commits.PR1655249
- Multiple rewrite rules aggregate Ethernet: (mpls-any and mpls-inet-both-non-vpn are NOT supported on PTX running BT ASIC). The order of applying the rewrite rules is not correct. The non-VPN rewrite rule is getting effect for the VPN traffic - See https://www.juniper.net/documentation/us/en/software/junos/cos/topics/conce pt/cos-rewriting-mpls-and-ipv4-packet-headers.htmlPR1655653
- In Junos OS Evolved PTX10001-36MR, LC1201 and LC1202, for a normal-mode scheduler (without strict-priority-scheduler in TCP), **priority medium-low** should map to hardware priority medium (same as **priority medium-high**).PR1656837
- Error logs from rpd/kernel corresponding to **JSR backup registration failed** might be observed during extreme scenarios of rpd restart etc. The issue is benign and self recovers with no impact expected on the connection or traffic. The right fix of the issue is in the latest primary branch but it cannot be backported as there risk associated with the fix complexity. PR1660685
- For ONDATRA test case which tested this with different network-instance names - default test is failing. For test to pass, use network-instance names - default.PR1662999
- GNOI API SetPackage through remote download is not supported. PR1665185
- On the PTX Junos OS Evolved platforms, MVRP enabled trunk ports with xSTP configured might go into blocked or designated state if the peer connects to the interface has no VLANs configured in its trunk port. PR1666921
- On a platform running Junos Evolved Software, SNMP GET for FPC MIBs might intermittently return "0" instead of the real value. PR1668285
- Default DDOS rate limit for LLDP packets is 20,000 PPS on PTX10003-60C, and on PTX10004 or PTX10008 or PTX10016 system with JNP10K-36QDD and JNP10K-4Q56DD-32Q28 line cards. PR1671196

- Firewall rule configuration of lo0 interface with option `next-header` on the egress direction is not supported. Configuring it could causes drop of other protocol traffic. PR1672315
- With **system configuration-database extend-size** feature enabled, the configuration validation fails during upgrade or downgrade or rollback to Junos OS Evolved 22.2R1 image, whereas if `no-validate` option is used during upgrade or downgrade or rollback to Junos OS Evolved 22.2R1 image (with **system configuration-database extend-size** feature enabled), then system goes to amnesiac mode. PR1672348

Interfaces and Chassis

- PTX10003: Sometimes 400G-ZR link does not come up when changed to channelized mode. PR1646915
- PTX10003: Link fails to come up on one of the channel in channelized mode after tuning. PR1656547

Infrastructure

- When using a source IP address as the management interface (with RPF check set to strict on interface), the response for the ICMP ping from the peer on the management interface is dropped by the Linux kernel as it expects the path to the peer to be on the WAN side. PR1498255
- Device can panic with vmcore under high memory pressure situations when kernel memory allocation fails on memdup_user(). PR1646610

Network Management and Monitoring

- When maximum-password-length is configured and user tries to configure password whose length exceeds configured maximum-password-length, error is thrown, along with error **ok** tag is also emitted. (Ideally **ok** tag should not be emitted in an error scenario.) The configuration does not get committed.PR1585855
- The mgd can crash when an invalid value is configured for identityref type leafs/leaf-lists while configuring Openconfig or any other third-party YANG, problem happens with json and xml loads. PR1615773

Routing Policy and Firewall Filters

- On Junos OS Evolved platforms, the unsupported configuration of BGP flow spec `interface-group exclude` might lead to errors and Packet Forwarding Engine corruption which does not permit filter bind. PR1639391

User Interface and Configuration

- Configuration archival (transfer-on-commit) over FTP protocol does not work in Junos OS Evolved. The reason is: internally a FreeBSD utility (fetch) is used to transfer the archived file. Junos OS Evolved is based on Linux and the fix or solution is to use a utility (cURL here) that is present in Linux.

```
Example config:
regress@vmore0r# show system archival
configuration {
    transfer-on-commit;
    archive-sites {
        "
ftp://regress:MaRtInI@sgrpca://var/tmp/"

 password "$9$rtneWX7NbY4J-VJUi.zFu0B"; ## SECRET-DATA
    }
}
```

PR1625937

- For an OpenConfig interface, if a matching interface-range is configured, then the deletion of that OpenConfig interface does not take effect. The workaround for this is - do not use interface-range configuration for OpenConfig interfaces. PR1637381

Resolved Issues

IN THIS SECTION

- Resolved Issues: 22.2R1 | **59**

Learn about the issues fixed in this release for PTX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

Resolved Issues: 22.2R1

IN THIS SECTION

- General Routing | 59
- Infrastructure | 63
- Interfaces and Chassis | 63
- MPLS | 63
- Network Management and Monitoring | 63
- Routing Protocols | 63
- User Interface and Configuration | 64

General Routing

- PTX10008: Junos OS Evolved, CB 1 becomes **Fault Standby** after `request node power-off re1`. PR1581476
- Filter with forwarding-class and destination-class combined might not work. PR1595788
- Junos OS Evolved: PTX series: An attacker sending a crafted GRE packet causes the PFE to restart (CVE-2022-22194). PR1614171
- the primary Routing Engine RE0 reloaded unexpectedly and new-primary Routing Engine RE1 does not bring up IS-IS or LDP adjacencies. PR1616114
- PTX Junos OS Evolved : **Output bytes** and **Output packets** counter values under Transit statistics on `show interfaces extensive` CLI decreases along with sending protocol control packets. PR1618587
- Wrong TX rate for queues configured **mix of high-low tx rates without excess bandwidth** on 10G interface might not work. PR1620284
- The cosd core file might be observed after Routing Engine switchover. PR1620758
- GNOI RPC TransferToRemote is not supported. PR1625212

- The primary role transfer might not be triggered on each rpd crashes if `switchover-on-routing-crash` is configured PR1625834
- Transient JSR replication errors 113 or 115 seen on disabling or enabling OSPF. PR1627625
- PTX10008 Junos OS Evolved : license installation fails with **validation hook evaluation failed** commit error. PR1628733
- cPTX: show system alarms shows many alarms which could be avoided. PR1628816
- PTX Junos OS Evolved : DDoS filter does not classify OSPF packets as OSPF-Hello and OSPF-Data packets. PR1628889
- Junos OS Evolved: Data streams two times during gNMI onchange streaming with hwdre/hwdfpc sensor paths. PR1632066
- P2MP LSP ping and trace from bud-node might fail when the branch is on another Packet Forwarding Engine. PR1632385
- [fabric] PTX10008: PDT: ERB : VxLAN: aggregated Ethernet lacp member link stuck is detached on PTX10008.PR1633849
- [vrrp] [vrrp_evo] PTX10004 :: Traffic loss is seen after VRRP primary role switchover on Junos OS Evolved 21.4. PR1633986
- Traffic impact might be seen when a firewall filter based policer for MPLS address family is configured on the device. PR1634644
- `show network-agent statistics detail` CLI output does not print expected output after sometime. PR1634716
- Label stack might be corrupted after Packet Forwarding Engine restart. PR1635130
- Client deadline might exceed with error after gribi route add with FIB ACK. PR1635727
- PTX10008 EVPN E-LAN: Ingress PE does not insert Sh label for BUM traffic received on local EP ESI interface, causing packet duplication on egress PE.PR1637703
- The rpd-agent crash might be observed once routing processes exit. PR1637391
- JTI UDP export support for /junos/system/cmerror/configuration and /junos/system/cmerror/counters does not work supported on chassis based systems like PTX10008 or PTX10016. PR1638262
- NPU util sensor to include FLT consumption for ZX and BT based PTX devices. PR1638487
- There is a mismatch between user-configured wavelength and actually transmitted wavelength on 400G-ZR wavelength setting with 75GHz spacing. PR1638603

- CCL:NGPR: RPD_KRT_RESPOSE_ERROR: krt change failed for prefix <> error from kernel is **EINVAL -- Bad parameter in request** PR1638745
- PTX10008 EVO : SNMP mib get on jnxLEDxxx generates general error with core file. PR1638768
- `show system errors error-id` CLI might show inconsistent threshold data.PR1640264
- Multicast packet drop might be seen when the outgoing interface flaps. PR1640294
- [Junos Telemetry Interface] Verification of DB data collection fails after executing the Junos Telemetry Interface decoder. PR1640442
- The `show network-agent statistics gnmi detail` CLI command reports packet drops for some gnmi target-defined mode sensors. PR1641483
- FPC start time is incorrect under `show chassis fpc details` CLI command. PR1641515
- [Telemetry] Filtering option for components name (CHASSIS, SIB) fails with /components/component sensor subscription. PR1641949
- Junos OS Evolved: In an MPLS scenario upon receipt of a specific IPv6 packet, an FPC crashes (CVE-2022-22214). PR1642721
- CFM traceoptions writes on every other line on PTX10008. PR1642948
- Traffic loss might be observed when you deactivate or activate the firewall filter. PR1643187
- The addition of new member to LAG might result in FPC crash. PR1643308
- jnxOperatingDRAMSize and jnxOperatingMemory OID values are incorrect on Junos OS Evolved platforms. PR1643910
- [fabric] : [[Junos OS Evolved-PTX10008] : PDT: ERB : VxLAN : Type5 traffic drop for BGP prefix on PTX10008 as remote leaf. PR1644458
- The interfaces might remain down and loopback wedge error might be seen. PR1645431
- The Routing Engine mastership might not transfer on each rpd crash. PR1645611
- Junos OS Evolved, disable operational commands. PR1646617
- CmInfra: Changing error severity does not work on FRUs managed by Routing Engine (like SIBs on PTX10008, FPCs on PTX1003 or TX10001-36MR.). PR1647282
- MAC learning might not happen on Junos OS Evolved PTX Series platforms. PR1647332
- PTX10003 is unable to forward traffic after the Layer 2 topology change. PR1647560

- High inter-packet delay and throughput performance degrade for Packet Forwarding Engine sensors. PR1648133
- Junos OS Evolved adding configuration `hash-key family inet layer-4` disables inet Hash-key Protocol. PR1648156
- Firewall counters might not increment for a longer time. PR1649324
- PTX10008: Junos OS Evolved : SyncE clock `hold-off-time` configuration does not work due to incorrectly computed timer value. PR1649358
- [CCL] debug-collector does not collect backup Routing Engine logs when SSH **root-login** configuration is set to **deny**. PR1649647
- HTTP or HTTPs as the tranfer for gnoi based remote file transfer is not supported. PR1650828
- The BFD session might flap in some scaled system with churn. PR1651473
- An error might be seen when the member link on an aggregated Ethernet bundle is deleted. PR1651932
- The rpd agent crash might be triggered after the interface flap for the backup Routing Engine. PR1652595
- P2MP traffic loss might be seen when link protected LSP reverts to the primary path. PR1652651
- DCF8: PTX10008: EVPN VXLAN intra-VLAN known unicast traffic flooded due to MAC installation failure on Packet Forwarding Engine. PR1652876
- PTX10008 Junos OS Evolved : `show snmp mib get` CLI returns incorrect value on jnxLED MIB OIDs. PR1654455
- The pkid core file is seen and can see interfaces lost. PR1655949
- The evo-aftmand-bt crash might be observed on Junos OS Evolved platforms. PR1657532
- The rpd might fail on backup Routing Engine on Junos OS Evolved platforms. PR1657797
- The packetio might generate core file when router reboot or FPC reboot is triggered. PR1658839
- The license might get out of sync between primary and backup Routing Engines. PR1658869
- PTX10008 Junos OS Evolved : hwdre core file is generated after Routing Engine switchover. PR1659377
- The BGP session might be flap on Junos OS Evolved platforms. PR1660805
- Channelized interface might go down if low-light-alarm or low-light-warning is enabled. PR1661215

- PTX10003 load balance v4_dscp and v6_dscp is enabled by default. PR1665131

Infrastructure

- The IPv6 default route might not take an effect in the global instance. PR1642576
- Configuring family MTU explicitly on an interface may cause host traffic to drop. PR1654140
- Traffic drop might be seen due to slow TCP reestablishment after a topology change. PR1661210

Interfaces and Chassis

- PTX10003 evo-aftmand process sees memory increasing linerally over days. PR1615000
- The snmp walk on jnxLEDTable fails on PTX10003. PR1620398
- Some daemons might get stuck when snmpd is at 100% CPU utilization. PR1636093
- [PTX10003] SSD DGM28-B56D81BCBQ || RE 0 SSD Primary minimum supported firmware version mismatch. PR1654762

MPLS

- IS-IS BFD sessions might take a long time to recover when the interface flaps. PR1593959
- LDP protections paths might not be established when `auto-targeted-session knob` is deactivated and activated. PR1620262

Network Management and Monitoring

- Configuring `set system no-hidden-commands` blocks or denies netconf or junoscript sessions. PR1590350
- False traffic spikes seen on SNMP graphs when ifHCOutOctets or ifHCInOctets are used. PR1635958
- The babeltrace core file might be triggered in a rare condition. PR1637992

Routing Protocols

- The BGP route might still be present in the multi-path route after increased IGP cost. PR1643665

User Interface and Configuration

- The addition or deletion of the gRPC configuration might cause a memory leak in the EDO app. PR1619974
- Passwordless authentication successful for configured user even after deleting ssh public key details from user login hierarchy. PR1625032
- [passive_monitoring] [monitoring] : PTX10008 :JDI_FT_REGRESSION: LAZURITE :: We observe configd object-info anomalies at net::juniper::config::interface::IFDCEtherOptionsCommon. PR1643192
- CCL:NGPR: configd core file during configd app restart test. PR1658688

Junos OS Evolved Release Notes for QFX5130-32CD, QFX5220, and QFX5700 Devices

IN THIS SECTION

- What's New | 65
- What's Changed | 71
- Known Limitations | 77
- Open Issues | 78
- Resolved Issues | 79

These release notes accompany Junos OS Evolved Release 22.2R1 for QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- What's New in 22.2R1 | 65

Learn about new features introduced in this release for QFX Series switches.

What's New in 22.2R1

IN THIS SECTION

- Class of Service | 65
- EVPN | 66
- Juniper Extension Toolkit (JET) | 67
- Junos Telemetry Interface | 68
- MPLS | 68
- OpenConfig | 69
- Routing Protocols | 69
- Additional Features | 70

Class of Service

- **Support for CoS within MPLS networks (QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use CoS within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. We have also added MPLS EXP rewrite support.

  - **Default CoS on the provider (P) and provider edge (PE) routers for MPLS interfaces**–The MPLS traffic uses the default EXP classifier. MPLS traffic is treated as best-effort traffic using the 802.1 default untrusted classifier. The default EXP classifier applies to all MPLS traffic on interfaces configured as `family mpls`. Differentiated Services code point (DSCP) classifiers are not applied to MPLS traffic.

- **Default CoS on PE routers for Layer 3 interfaces**–By default, all Layer 3 VPN logical interfaces are bound to default DSCP classifiers.

  If you apply an EXP classifier on a penultimate-hop popping (PHP) node, then by default, the MPLS header TLL value overwrites the IP header time-to-live (TTL) value. In this case, a zero (0) overwrites the IP header DSCP bits, which signifies uniform mode. To use pipe mode, where nothing overwrites IP header TTL values or IP header DSCP bits, you should configure the following command:

  ```
  set protocols mpls no-propagate-ttl
  ```

  NOTE: The DSCP of IP in MPLS packets can't be remarked either at PE or P routers.

  [See Understanding CoS MPLS EXP Classifiers and Rewrite Rules.]

EVPN

- **Assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN edge-routed bridging (ERB) fabric (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can configure AR and OISM together in an EVPN-VXLAN ERB overlay fabric. These devices can serve as AR replicator devices or AR leaf devices.

  Here is a summary of integrated AR and OISM support:

  - AR replicator devices are lean spine devices in the fabric. The AR replicator function can't be collocated with the OISM border leaf or server leaf role on the same device.
  - AR leaf devices can be OISM server leaf or border leaf devices.
  - AR replicator devices must be running Junos OS Evolved software that supports OISM (even though the AR replicator devices don't operate in either an OISM server leaf or OISM border leaf role).

  When you configure AR devices:

  - You configure the EVPN instances using MAC-VRF instances with `vlan-based` or `vlan-aware` service types only.
  - You must configure the AR replicator devices with the same tenant VRF instances, corresponding IRB interfaces, and member VLANs as the OISM devices.

  [See Assisted Replication Multicast Optimization in EVPN Networks and Optimized Inter-Subnet Multicast in EVPN Networks.]

- **Blocking asymmetric EVPN Type 5 routes (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can configure the local node to reject asymmetric EVPN Type 5 routes on EVPN-VXLAN networks. The local node examines the incoming EVPN Type 5 route packets and

rejects the route when the virtual network identifier (VNI) in the ingress route differs from the locally configured VNI.

To block asymmetric EVPN Type 5 routes, include the `reject-asymmetric-vni` statement at the `[edit routing-instance` *`routing-instance-name`* `protocols evpn ip-prefix-routes]` hierarchy level.

[See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN and ip-prefix-routes.]

- **Interconnect EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved 22.2R1, we support Data Center Interconnect (DCI) stitching for EVPN-VXLAN gateway tunnels. The gateway connects the data center and the WAN, and both data center and WAN gain forwarding states for route distinguishing, route targeting, and interconnect Ethernet segment identifier (I-ESI) support.

  DCI control plane stitching also:

  - Supports multihoming.
  - Extends the Layer 2 (L2) connectivity required for some tenants in a data center.
  - Uses the unknown MAC route to prevent MAC scale issues on data center network virtualization edge (NVE) devices.

  [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.]

- **Symmetric IRB with EVPN Type 2 routes (ACX7100, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)**—Starting in Junos OS Evolved Release 22.2R1, you can enable symmetric IRB EVPN Type 2 routing in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) edge-routed bridging (ERB) overlay fabric. With the symmetric routing model, leaf devices can route and bridge traffic on both ingress and egress sides of a VXLAN tunnel. Leaf devices use a transit VXLAN network identifier (VNI) and Layer 3 (L3) interfaces on the associated VLAN to exchange traffic across the VXLAN tunnels.

  We support this feature with `vlan-aware` and `vlan-based` MAC-VRF instance service type configurations. To enable this feature, you must also configure EVPN Type 5 routing with L3 VRF instances to establish intersubnet reachability among the EVPN devices.

  [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics and irb-symmetric-routing.]

Juniper Extension Toolkit (JET)

- **JET API client gRPC connections limit (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, we have increased the maximum number of Juniper Extension Toolkit (JET) gRPC connections to 300. Use the

`max-connections` statement at the `[edit system services extension-service request-response grpc]` hierarchy level to configure the maximum number of gRPC connections.

[See max-connections.]

Junos Telemetry Interface

- **OpenConfig telemetry system model support (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130, and QFX5220)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig data models openconfig-telemetry.yang and openconfig-telemetry-types.yang. Support includes streaming of state data for dynamic subscriptions and configuration for persistent subscriptions.

  [See Telemetry Sensor Explorer for telemetry support and OpenConfig User Guide for configuration.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

MPLS

- **Enhancements to BFD-triggered MPLS fast reroute (FRR) for unicast next hops and session-id-change-limiter-indirect (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700, and QFX5200)**—In Junos OS Evolved Release 22.2R1, we've enhanced the BFD-triggered fast reroute for unicast next hops and session-id-change-limiter-indirect to address the issue of traffic being silently discarded because of a session mismatch between the control plane and data plane.

  You can limit the re-programming of the number of parent nodes of the indirect-nexthop and prevent additional complexity in the Packet Forwarding Engine when the session-identifier id of the indirect nexthop is changed. To make those changes, use the `session-id-change-limiter-indirect` configuration statement at the `[edit routing-options]` hierarchy level.

  [See Bidirectional Forwarding Detection (BFD) for MPLS and session-id-change-limiter-indirect.]

OpenConfig

- **OpenConfig telemetry system model support (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130, and QFX5220)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig data models openconfig-telemetry.yang and openconfig-telemetry-types.yang. Support includes streaming of state data for dynamic subscriptions and configuration for persistent subscriptions.

  [See Telemetry Sensor Explorer for telemetry support and Mapping OpenConfig Telemetry System Model Commands to Junos Configuration for configuration.]

- **VLAN data model support (ACX7100, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016 and QFX5130)**—Junos OS Evolved Release 22.2R1 introduces support for OpenConfig VLAN configuration based on the following data models:

  - **openconfig-vlan.yang, version 3.2.1**
  - **openconfig-network-instance.yang, version 0.16.1**
  - **openconfig-network-instance-l2.yang, version 0.16.1**

  Support includes OpenConfig configuration and streaming of state data.

  [See Telemetry Sensor Explorer for state sensors and Mapping OpenConfig VLAN Commands to Junos Configuration for configuration.]

- You configure the JSD port with the command `set services extension-service request-response grpc clear-text port` *`port-number`* `max-connections` *`number`*. You use OpenConfig to configure that port as a p4rt port. You then create a Layer 2 and Layer 3 interfaces configuration on the device with VLAN tags and the IP address for the device, allowing Layer 2 packets to be punted and injected between the device and the SDN controller. You use firewall filters to direct packet punts to the p4-switch.

  [See OpenConfig User Guide.]

Routing Protocols

- **TCP authentication option (TCP-AO) for resource public key infrastructure (RPKI) validation sessions (ACX7100-32C, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, you can use the TCP authentication option to authenticate RPKI validation sessions for securing the Internet's routing infrastructure, such as the BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

  To enable a TCP authentication option chain to authenticate an RPKI validation session, use the configured `authentication-algorithm` **`ao`** `and authentication-key-chain` *`keychain`* at the [`edit routing-options validation group` *`group_name`* `session` *`address`*] and [`edit routing-options validation group` *`group_name`*] hierarchy level.

[See TCP Authentication Option (TCP-AO)]

- **BGP extended route retention (PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 22.2R1, we have enhanced the long-lived graceful restart (LLGR) capabilities for BGP Helper and introduced the new configuration statement `extended-route-retention` at `[edit protocols bgp group neighbor graceful-restart long-lived]` hierarchy level. When you enable this feature, `extended-route-retention` supports LLGR helper mode regardless of the BGP peer LLGR capabilities. We have also updated the output of the following existing operational commands:

  - show bgp neighbor
  - show route extensive

  [See graceful-restart-long-lived-edit-protocols-bgp.]

- **RPD Object Reference count Anomaly checker (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5220)**— Starting in Junos OS Evolved Release 22.2R1, we introduce a generic reference count infrastructure that all the modules in the routing protocol process (rpd) can use. The module maintains lock and unlock statistics corresponding to each object type in use. Any application can call the refcount increment or decrement API when an object is referred or deferenced. The module also provides a mechanism to detect anomalies such as a leak or overflow in an object's refcount.

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **Collect ON_CHANGE BGP RIB telemetry statistics** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5220)

  [See Telemetry Sensor Explorer.]

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding** (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, and QFX5700)

  [See Telemetry Sensor Explorer.]

- **Media Access Control Security (MACsec)** (QFX5700 with 20x25G line cards)—

  Support includes:

  - AES-128 and AES-256 encryption
  - Extended packet numbering
  - Fallback preshared key (PSK)

- Timer-based MACsec secure association key (SAK) refresh
- Jumbo MTU support
- Fail-open mode
- PSK hitless rollover

[See Understanding Media Access Control Security (MACsec).]

- **Supported transceivers, optical interfaces, and direct attach copper (DAC) cables** —Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update this tool and provide the first supported release information when the optic becomes available.
- **Support for TCP authentication and keychain** (QFX5700 and QFX5220)—We asupport TCP authentication and key update mechanism for BGP and LDP routing protocols, which enables you to update authentication keys without interrupting associated routing and signaling protocols.

  [See Authentication for Routing Protocols, authentication-key-chain (Protocols LDP), authentication-algorithm, and authentication-key-chain (Protocols BGP and BMP).]
- **Support for TCP Authentication Option (TCP-AO)** (QFX5220)—You can now use TCP-AO to authenticate TCP segments exchanged during BGP and LDP sessions.

  [See TCP Authentication Option (TCP-AO) and authentication-key-chains (TCP-AO).]

What's Changed

IN THIS SECTION

- What's Changed in Release 22.2R1-S1 | **72**
- What's Changed in Release 22.2R1 | **72**

Learn about what changed in these releases for QFX Series switches.

What's Changed in Release 22.2R1-S1

IN THIS SECTION

- Juniper Extension Toolkit (JET) | 72
- Junos OS API and Scripting | 72

Juniper Extension Toolkit (JET)

- **Python 3 is the default and only Python version for executing Juniper Extension Toolkit Python scripts (ACX Series, PTX Series, and QFX Series)**—Junos OS Evolved supports only Python 3 for executing Juniper Extension Toolkit (JET) scripts written in Python. Python 2.7 is no longer supported for executing JET scripts, and we've deprecated the `language python` statement at the `[edit system scripts]` hierarchy level.

  [See Understanding Python Automation Scripts for Junos Devices.]

Junos OS API and Scripting

- **Deprecated functions in the `libpyvrf` Python module (ACX Series, PTX Series, and QFX Series)**—The `libpyvrf`Python module no longer supports the `get_task_vrf()` and `set_task_vrf()` functions.

  [See How to Specify the Routing Instance in Python 3 Applications on Devices Running Junos OS Evolved.]

What's Changed in Release 22.2R1

IN THIS SECTION

- General Routing | 73
- Interfaces and Chassis | 73
- OpenConfig | 73
- Multicast | 74
- Network Management and Monitoring | 74
- Routing Protocols | 75
- Timing and Synchronization | 75

User Interface and Configuration | 75

General Routing

- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series and QFX Series)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route `set routing-options?`. Any instance you explicitly configure is translated into `set routing-instance r1 routing-options?`. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.

- Stateful port configuration for PTP over Ethernet and default profile is supported only on boundary clock mode and not on ordinary clock mode.

- **DSCP rewrite on EVPN VXLAN NNI ports (QFX5130 and QFX5700)**— QFX5130 and QFX5700 platforms support DSCP rewrite on EVPN VXLAN NNI ports with limitations.

  [See Implementing CoS on VXLAN Interfaces (Junos OS Evolved.]

- **Change in in unnumbered-address support for GRE tunnel**—There is a behavioural change in unnumbered-address support for GRE tunnel with IPv6 family and display donor interface for both IPv4 and IPv6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level.

  [See show interfaces.]

Interfaces and Chassis

- **Display the donor details of the IPv6 borrower interface**— The output for the show interfaces command now displays the donor details of the IPv6 borrower interface.

  [See show interfaces.]

OpenConfig

- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".

Multicast

- **Changes to `show mvpn c-multicast` and `show mvpn instance` outputs**— The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

  [See show mvpn c-multicast.]

  .]

Network Management and Monitoring

- **Changes to the NETCONF `<edit-config>` RPC response (ACX Series, PTX Series, and QFX Series)**— When the `<edit-config>` operation returns an error, the NETCONF server does not emit a `<load-error-count>` element in the RPC response. In earlier releases, the `<edit-config>` RPC response includes the `<load-error-count>` element when the operation fails.

- **Enhanced system log messages (QFX5130, QFX5200, QFX5220, and QFX5700)**— We've added multiple events inside the event tag using the **UI_LOGIN_EVENT|UI_LOGOUT_EVENT** format, which has an option (|) to separate the events, to generate system log messages.

  Earlier to this release, the event tag used the **UI_LOGIN_EVENT UI_LOGOUT_EVENT** format and for various combinations of rpc filters was not getting logged.

  [See Overview of System Logging https://www.juniper.net/documentation/us/en/software/junos/network-mgmt/topics/topic-map/system-logging.html#id-overview-of-junos-os-system-log-messages.]

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**— The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

  [See max-datasize.]

- **Change in behavior of SNMP MIB object ifAlias**— SNMP MIB object ifAlias now shows the configured interface alias. In earlier releases, ifAlias used to show configured interface description.

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

  See privacy-3des and privacy-aes128

Routing Protocols

- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the `[edit system services ssh]` hierarchy level.

  In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the `[edit system services ssh]` hierarchy level.

- **The RPD_OSPF_LDP_SYNC message not logged**—On all Junos OS and Junos OS Evolved devices, when LDP session goes down there is a loss of synchronization between LDP and OSPF. After this loss of synchronization, when an interface is in the holddown state for more than 3 minutes, a system log message with a warning level is sent. This message appears in both the messages file and the trace file.

  However, this system log message does not get logged when you have explicitly configured the `hold-time` for `ldp-synchronization` at the `[edit protocols ospf area` *`area id`* `interface` *`interface name`* hierarchy level . If the configured `ldp-synchronization hold-time` is less than 3 minutes, then the syslog is not printed. However, the message is printed after 3 minutes.

- To achieve consistency among resource paths, the resource path/**/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter ip-addr='address'/state/countersname='name'/out-pkts/** is changed to /mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'**/state/counters name='name'/**. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

- When the krt-nexthop-ack statement is configured, the RPD will wait for the next hop to get acknowledged by PFE before using it for a route. Currently, only BGP-labeled routes and RSVP routes support this statement. All other routes will ignore this statement.

Timing and Synchronization

- **Performance monitoring time interval with UTC on Junos OS Evolved platforms** —The performance monitoring (PM) time interval for 1-day bins on Junos OS Evolved platforms begins at midnight in the UTC zone, aligning with the standard behaviour of Junos OS. This synchronization allows you to maintain consistent performance monitoring schedules across platforms, enhancing the accuracy and reliability of network performance data.

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new

CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.

- **Load JSON configuration data with unordered list entries (ACX Series, PTX Series, and QFX Series)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:

  - Use the `request system convert-json-configuration` operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.

  - Configure the `reorder-list-keys` statement at the `[edit system configuration input format json]` hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.

  When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

  [See json and request system convert-json-configuration.]

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.

- **Persistent CLI timestamps**—To have a persistent CLI timestamp for the user currently logged in, enable the `set cli timestamp` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements:

```
set system login class class name permissions permissions
set system login class class name cli timestamp set system login user username class class
name authentication plain-text-password
```

NOTE: The default timestamp format is %b %d %T. You can modify the format per your requirements. For example, you can configure the following statement:

`set system login class class name cli timestamp format "%T %b %d`

To enable timestamp for a particular user with default class permissions and format, configure the following statements:

```
set system login user username class class name authentication plain-text-password
set system login user username cli timestamp
```

Known Limitations

IN THIS SECTION

- General Routing | 77

Learn about known limitations in Junos OS Evolved Release 22.2R1 for QFX Series switches.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- On QFX5220-32CD, VLANs between 3968 and 4095 are reserved for Layer 3 interfaces by default. So, these VLANs cannot be used for Layer 2 interfaces. As of now there is no commit check added for this purpose. You need to take care of this while configuring VLANs for Layer 2. PR1423468
- On QFX5700 when the FEB card has a power failure the `show system alarms` CLI displays the following error **PCI Uncorrected error on dev 0000:00:03.2** along with FEB alarm. PR1578066

Open Issues

IN THIS SECTION

- General Routing | 78

Learn about open issues in Junos OS Evolved Release 22.2R1 for QFX Series switches.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- On QFX5700 platforms few interfaces do not come up after removing channelization through single commit, that is by using `delete interfaces`. PR1592238
- On QFX5700 ungraceful removal (OIR) of FPC or an FPC fault might result in a PCIE MAJOR alarm **PCI Uncorrected error on dev 0000:00:03.0** which does not get cleared. The only way to clear this alarm is reboot of the device. There are 2 situations in which this alarm can be seen:

  1. FPC is faulty: In rare FPC fault cases, the PCI Uncorrected error alarm might be seen along with FPC going to a Fault state as indicated by the `show chassis fpc` command. This is accompanied by other FPC major alarms. Once the faulty FPC is replaced with a good one, the alarm is still seen, and a reboot is required to clear this alarm. Post identification of the fault and FPC replacement, this alarm is harmless, and FPC state can be confirmed by issuing the `show chassis fpc` command.
  2. Ungraceful OIR: The ungraceful removal of FPCs is not recommended on QFX5700. This operation might result in PCI Uncorrected Error alarm. Use one of the following two methods to do a graceful FPC OIR removal:

     a. Execute the `request chassis fpc slot <slot #> offline` command from the CLI.

     b. Press the Offline Button for 1 second on the FPC to offline the FPC.

     Once the FPC is gracefully offlined both LEDs - PWR and STS will go off. The FPC can be removed at this point.

     PR1620197

- 400G LR4-10 link does not come up after deleting interface disable configuration when port is disabled followed by system reboot. PR1625494
- Interface convergence time for QFX5220-128C is 1 min more than earlier. PR1636181
- Unexpected carrier transitions are seen on JNP-100G-2X50G-xM after cable plug out and plug in on QFX5130-32CD and QFX5220-32CD. It is observed that the Carrier Transitions counter increments by 1 on cable plug out and by 3 on cable plug in. It should increment by 1 in both cases. PR1642744
- During bulk mac move scenarios, we see few MAC entries missing in the Routing Engine CLI `show ethernet-switching table`, even though those entries are present in the hardware. As MAC is present in the hardware, there is no impact to the Packet Forwarding Engine. PR1650329
- TOS(DSCP+ECN) bits are not getting copied from the Inner Layer 3 header to outer VXLAN header at the Ingress VTEP. Because of this in the core, ECN marking and DSCP classification are not working.PR1658142
- On QFX5700, if JNP-FPC-20Y is restarted, in rare situations traffic flow does not resume on all ports of the FPC.PR1659566
- The general recommendation with OISM is to not have CE interfaces in SBD. Adding the same can lead to traffic drops. PR1668921
- After evo-pfemand restart, one interface does not come up. PR1669037
- On QFX5130-32CD, ND resolve packets are subjected to loopback filters configured. The resolve packets might get dropped as per the loopback filter configuration. In some cases, if IPv6 neighbours are cleared using `clear ipv6 neighbors` command, the device may not re-learn those IPv6 neighbours. PR1671730

Resolved Issues

IN THIS SECTION

- Resolved Issues: 22.2R1 | **80**

Learn about the issues fixed in this release for QFX Series switches.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

Resolved Issues: 22.2R1

IN THIS SECTION

- General Routing | 80
- Infrastructure | 80
- Network Management and Monitoring | 80
- Routing Protocols | 81

General Routing

- QFX10003 Layer 2: In scaled Layer 2 network, error logs are printed for MAC Creation. MAC-learning works as expected. PR1491933
- Filter with forwarding-class and destination-class combined might not work. PR1595788
- After sigkill/app crash, jstatsd app does not come up. PR1641229
- FEC corrected errors which are cleared with clear statistics command show up as huge value after unified ISSU. PR1641583
- A forced reboot might be observed when SSD is not detected during a script or daemon call. PR1648117
- Junos OS Evolved adding configuration `hash-key family inet layer-4` disables inet Hash-key Protocol. PR1648156
- Junos OS Evolved, QFX **EvoPfemand-main** process memory leak. PR1652873

Infrastructure

- On QFX5220 devices, change in the output for L3VPN.inet6 route table in `show route forwarding-table summary | display xml` command. PR1653182

Network Management and Monitoring

- False traffic spikes seen SNMP graphs when ifHCOutOctets or ifHCInOctets are used. PR1635958

Routing Protocols

Junos OS Evolved: Denial of Service (DoS) vulnerability in RPD upon receipt of specific BGP update (CVE-2022-22213). PR1642741

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700.

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: https://www.juniper.net/support/downloads/
2. In the Find a Product box, enter the Junos OS platform for the software that you want to download.
3. Select Junos OS Evolved from the OS drop-down list.
4. Select the relevant release number from the Version drop-down list.
5. In the **Install Package** section, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

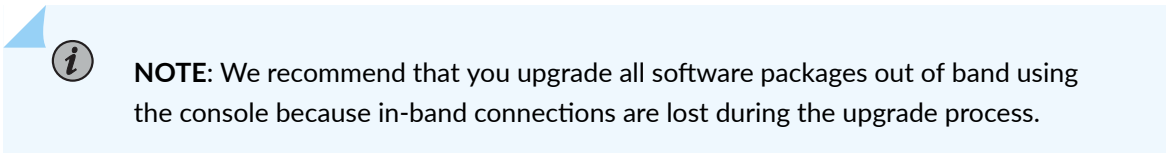

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

For more information about software installation and upgrade, see Software Installation and Upgrade Overview (Junos OS Evolved). For more information about EOL releases and to review a list of EOL releases, see https://support.juniper.net/support/eol/software/junosevo/.

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see Juniper Flex Program.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

  https://apps.juniper.net/feature-explorer/

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

  https://prsearch.juniper.net/InfoCenter/index?page=prsearch

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

  https://apps.juniper.net/hct/home

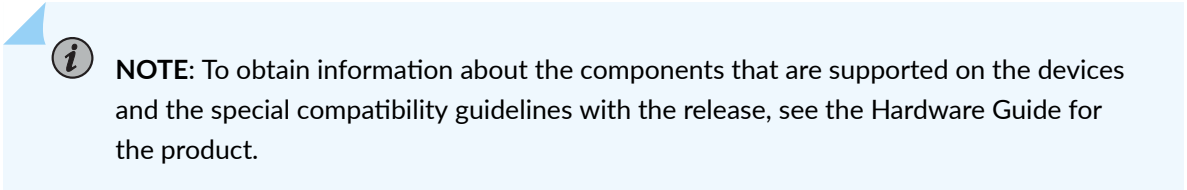

NOTE: To obtain information about the components that are supported on the devices and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6.

  https://pathfinder.juniper.net/compliance/

Requesting Technical Support

IN THIS SECTION

- Self-Help Online Tools and Resources | 83
- Creating a Service Request with JTAC | 84

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.pdf.
- Product warranties—For product warranty information, visit https://support.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://support.juniper.net/support/
- Search for known bugs: https://prsearch.juniper.net/

- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://supportportal.juniper.net/s/knowledge
- Download the latest versions of software and review release notes: https://support.juniper.net/support/downloads/
- Search technical bulletins for relevant hardware and software notifications: https://supportportal.juniper.net/s/knowledge
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://supportportal.juniper.net/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://support.juniper.net/support/requesting-support/
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

Revision History

28 March 2025—Revision 13, Junos OS Evolved Release 22.2R1

27 March 2024—Revision 12, Junos OS Evolved Release 22.2R1

7 March 2024—Revision 11, Junos OS Evolved Release 22.2R1

10 August 2023—Revision 10, Junos OS Evolved Release 22.2R1

20 July 2023—Revision 9, Junos OS Evolved Release 22.2R1

20 April 2023—Revision 8, Junos OS Evolved Release 22.2R1

6 April 2023—Revision 7, Junos OS Evolved Release 22.2R1

24 November 2022—Revision 6, Junos OS Evolved Release 22.2R1

22 August 2022—Revision 5, Junos OS Evolved Release 22.2R1

8 August 2022—Revision 4, Junos OS Evolved Release 22.2R1

4 August 2022—Revision 3, Junos OS Evolved Release 22.2R1

7 July 2022—Revision 2, Junos OS Evolved Release 22.2R1

23 June 2022—Revision 1, Junos OS Evolved Release 22.2R1