Release Notes

Published
2026-04-13

Juniper Routing Director Release 2.8.0

Software Highlights

Device Life-Cycle Management

- Device Support:
  - Basic device support (onboard devices, export device details, view audit logs, and upgrade software)—SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, and EX4400 Virtual Chassis.
  - Full support (device observability, routing observability, service orchestration, active assurance, network optimization, and viewing topology weather map)—MX301.
- Use management VRF to adopt brownfield Junos devices.
- Support for gNMI dial-in connection for telemetry.

Observability

- AI/ML-driven device health monitoring on additional platforms (PTX10004 and PTX10008).
- Detect physical layer faults in Routing Director using AI/ML.
- Increase KPI data retention.
- Use Export Manager to stream data to external systems.
- IGP anomaly alert generation.
- View BGP route details.

Service Orchestration

- Updated ESI configuration format in the L2 resource instance.
- Configure custom local VPWS ID.

Network Optimization

- Re-parse device collection.
- Map display options on the Topology page.

- Flexible metric selection for color legend.
- Configure multiple autonomous systems per BGP-LS.

Planner

- View offline topology changes.
- Group working models using tags.
- View reports with in the Topology page.

Active Assurance

- Centralized view of Measurement reports.
- Support for additional plug-ins (Cisco IP SLA Ping, Path MTU Discovery, and QoS Policy Profiling).

Installation and Upgrade

- Install Juniper Routing Director on a single node.

Table of Contents

Introduction | 1

Licensing | 2

Supported Junos OS Releases, Devices, and Browsers | 3

New Features | 10

Known Issues | 20

Resolved Issues | 39

Introduction

Service providers, cloud providers, and enterprises are facing an increase in the volume, velocity, and types of traffic. This creates both unique challenges (increased user expectations and expanded security threats) and fresh opportunities (new generation of 5G, IoT, distributed edge services) for network operators.

To accommodate rapid changes in traffic patterns, service providers and enterprises need to quickly detect and troubleshoot devices and service issues, and make changes to service configurations in real-time. Any misconfiguration due to human errors can lead to service outages. Investigating and resolving these issues can be a time-consuming process.

Juniper®Routing Director is a WAN automation solution that enables service provider and enterprise networks to meet these challenges. Juniper's solution delivers an experience-first and automation-driven network that provides a high-quality experience to network operators.

Routing Director is based on a modern microservices architecture with open APIs. Routing Director is designed with an easy to use UI that provides a superior operational and user experience. For example, Routing Director implements different persona profiles (such as network architect, network planner, field technician, and Network Operations Center [NOC] engineer) to enable operators to understand and perform the different activities in the device life-cycle management (LCM) process.

Routing Director takes a use case-based approach to network operations. When you execute a use case, Routing Director invokes all the required capabilities of that use case, runs a workflow (if necessary) and presents you with a completed set of tasks that implements the use case.

Routing Director supports the following use cases:

- **Device life-cycle management (LCM)**—Allows you to onboard, provision, and then manage a device. Routing Director automates the device onboarding experience, from shipment through service provisioning, thus enabling the device to be ready to accept production traffic.
- **Observability**—Allows you to visualize the network topology, provision tunnels, view topology updates in real-time, and monitor devices and the network. You can also view device and network health and drill down into the details. In addition, Routing Director notifies you about network issues using alerts, alarms and events, which you can use to troubleshoot issues affecting your network. Routing Director also provides a routing dashboard and an interactive routing topology map where you can actively monitor the overall routing health of your network in real time.
- **Trust and compliance**—Automatically checks whether the device complies with the rules defined in the Center for Internet Security (CIS) benchmarks document. In addition, Routing Director also checks the configuration, integrity, and performance of the device and then generates a trust score that determines the device's trustworthiness.

- **Service Orchestration**—Enables you to streamline and optimize the delivery of network services, thereby improving efficiency and reducing the risk of errors. A service can be any point-to-point, point-to-multipoint or multipoint-to-multipoint connection. For example, Layer 3 VPNs or EVPNs.
- **Active Assurance**—Enables you to actively monitor and test the network's data plane by generating synthetic traffic using Test Agents. Test Agents are measurement points deployed in certain routers in your network. These Test Agents are capable of generating, receiving, and analyzing network traffic and therefore enable you to continuously view and monitor both real-time and aggregated result metrics.
- **Network Optimization**—Enables you to optimize the utilization of network resources, enhance network performance, and ensure reliable and efficient delivery of data across the network. Routing Director optimizes the network by managing the life-cycle of label-switched paths (LSPs) or segment routing policies, through an intent-based approach.
- **Network Planner**— Provides in-depth network views and reports on how the network is performing in a particular failure scenario, all without impacting your production environment.

For details about these use cases and other features of Routing Director, see "New Features" on page 10.

In summary, Routing Director helps operators to automate the onboarding and provisioning of devices, simplify and accelerate service delivery, evaluate device and service performance, and reduce manual effort and timelines.

Use these release notes to know about features, supported Junos OS and Junos OS Evolved releases, supported devices, and open issues in Routing Director.

Licensing

To use Routing Director and its features, you need:

- **Product Entitlement**—To use Routing Director and its use cases.

  NOTE: Product entitlements are honor-based and not enforced for Routing Director Release 2.8.0.

- **Device License**—To use the features on a device that you onboarded.

To purchase a license, contact your Juniper Networks sales representative. For more information about purchasing licenses, see Juniper Licensing User Guide. After you purchase a license, you can download the license file and manage licenses by using the Juniper Agile Licensing (JAL) portal. You can also

choose to receive the license file over an e-mail. The license file contains the license key. The license key determines whether you are eligible to use the licensed features.

After the device is onboarded, the Super User and the Network Admin can add a device license from the **Licenses** tab (**Observability** > **Health** > **Troubleshoot Devices** > ***Device-Name*** > **Inventory** > **Licenses**) of the Routing Director GUI. For more information, see Manage Device Licenses.

Supported Junos OS Releases, Devices, and Browsers

IN THIS SECTION

- Supported Juniper Networks Devices | 4
- Supported Non-Juniper Devices | 6
- Supported Devices for Test Agents | 8
- Supported OS for RFC2544 | 8
- gNMI Support | 9
- Supported Browsers | 9

This topic includes information on:

- Devices (Juniper Networks and third-party) supported by Routing Director
- Supported operating system (OS) versions
- Supported devices for Test Agents
- Supported OS for RFC2544
- gNMI support for Juniper Networks devices
- Supported browsers

Supported Juniper Networks Devices

Table 1 on page 4 lists all the devices supported by Routing Director and the supported operating system (OS) versions.

Table 1: Supported Juniper Devices and OS Versions in Routing Director

Device Family	Device Series	Supported OS Version
ACX Series	• ACX710 (Only device management functions, custom rules and custom service designs) • ACX2200 (EMS functionality and topology-related information only) • ACX5048, ACX5096 (Only device management functions, custom rules, and custom service designs) • ACX5448 (Only device management functions, custom rules, and custom service designs) • ACX6360 (device management functions) • ACX 7020, ACX7024, ACX7024-X • ACX7100-32C, ACX7100-48L • ACX7332 • ACX7348 • ACX7509	Junos OS Evolved releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R1, 24.2R2, and 24.4R1
PTX Series	PTX10001-36MR, PTX10002-36QDO, PTX10004, PTX10008, PTX10016	Junos OS Evolved releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R2, and 24.4R1 Junos OS 22.4R2

Table 1: Supported Juniper Devices and OS Versions in Routing Director *(Continued)*

Device Family	Device Series	Supported OS Version
MX Series	• MX104 (device management and observability functions only) • MX301 • MX204, MX240, MX304, MX480, MX960 • MX2008, MX2010, MX2020 • MX10003, MX10004, MX10008, MX10016 • vMX	Junos OS Releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R2, and 24.4R1
EX Series	• EX3400 • EX2300 (Only device management functions, custom rules and custom service designs) • EX3400 • EX4300 [EX4300-32F (EMS only), EX4300-48mp] • EX4400 with Virtual Chassis (Only device management functions). • EX9200 **NOTE**: Network management functions and telemetry are supported only on EX4300-48MP. Support on other EX4300 models is limited to basic network management functions such as device reboot, software upgrade, configuration backup, audit logs, and so on.	Junos OS releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R2, and 24.4R1

Table 1: Supported Juniper Devices and OS Versions in Routing Director *(Continued)*

Device Family	Device Series	Supported OS Version
QFX Series	• QFX5110, QFX5120, QFX5200	Junos OS releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R2, and 24.4R1
SRX Series Firewalls	• SRX5400, SRX5600, SRX5800 (Only device management functions, custom rules and custom service designs) • SRX4100, SRX4200, SRX4300, SRX4600, and SRX4700 (Only onboard devices, export device details, view audit logs, and upgrade software)	Junos OS releases 22.2R3, 22.4R2, 23.2R2, 23.4R2, 24.2R2, and 24.4R1
vMx	-	• Junos OS releases 24.2R2, 23.4R2, 23.2R2, 22.4R2, and 22.2R3.

Supported Non-Juniper Devices

Table 2 on page 7 lists the non-Juniper devices supported by Routing Director.

NOTE: For third-party devices:

- You can perform basic device management functions (such as, basic device adoption, execute simple gNOI commands (reboot), and add configuration templates) and onboard devices using both GUI and API.
- You cannot enable routing protocol analytics.
- You can collect telemetry data for Cisco devices by deploying custom rules. However, this is a beta feature.
- You can configure PCEP for Cisco IOS XRv . We support pce-initiated, device-controlled and delegated tunnels.

- Compared to Juniper Networks and Cisco devices, the support for Nokia devices remains limited.

Table 2: Supported Non-Juniper Devices and OS Versions in Routing Director

Vendor	Device Model	OS Version
Nokia	• Nokia 7750 SR-s • Nokia 7750 SR-12e • Nokia 7750 SR-a8 • Nokia 7250 IXR-x • Nokia 7250 IXR-e	TiMOS-C-25.3.R2
Cisco Systems	Cisco 8201	IOS-XR 7.9.0 and IOS-XR 24.3.2
	Cisco 8202	IOS-XR 7.9.0 and IOS-XR 24.3.2
	Cisco NCS 57C3	IOS-XR 24.3.2
	Cisco NCS 5504	IOS-XR 7.3.2 and IOS-XR 24.3.2
	Cisco NCS 5508	IOS-XR 7.3.2
	Cisco IOS XRv	IOS-XR 7.0.0 and IOS-XR 24.3.1
	Cisco ASR 9902	IOS-XR 7.11.2
	Cisco ASR 9900	IOS-XR 7.11.2
	Cisco ASR 9910	IOS-XR 7.11.2

Supported Devices for Test Agents

Table 1 on page 4 lists all the devices supported for Test Agents.

Table 3: Supported Devices and OS Versions for Test Agents

Device Family	Device	Minimum Supported Release
ACX Series	ACX7100-48L and ACX7100-32C	Junos OS Evolved Release 23.1
	ACX7509	Junos OS Evolved Release 23.1
	ACX7020-AC and ACX7020-DC	Junos OS Evolved Release 24.4
	ACX7024	Junos OS Evolved Release 23.1
	ACX7024X	Junos OS Evolved Release 24.1
	ACX7348	Junos OS Evolved Release 23.4
	ACX7332	Junos OS Evolved Release 23.4
PTX Series	PTX10001-36MR, PTX10002, PTX10002-36QDD, PTX10008	Junos OS Evolved Release 22.3
MX Series	MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10004, MX10008, vMX	Junos OS Release 22.3

Supported OS for RFC2544

Table 4 on page 9 lists the supported operating systems (OS) for RFC2544.

Table 4: Supported OS for RFC2544

Plug-in	Junos OS	Junos OS Evolved
RPM RFC2544 CCC RPM RFC2544 ETH RPM RFC2544 INET	Not supported	Supported. Only Release 23.4 is supported.
RPM RFC2544 CCC reflector RPM RFC2544 ETH reflector, RPM RFC2544 INET reflector	Supported	Supported

gNMI Support

Table 5 on page 9 lists the gNMI support for Juniper Networks devices.

Table 5: gNMI Support on Juniper Networks Devices

Device Series and Model	gNMI Support
ACX Series	Starting from Junos OS Evolved Release 22.4 and later. **NOTE**: ACX2200, ACX5048, and ACX5096 devices do not support gNMI.
PTX Series	Starting from Junos OS Evolved Release 22.4 and later.
MX Series	Starting from Junos OS Release 22.4R1 and later. **NOTE**: MX104 devices do not support gNMI.
EX Series	Supported only on EX4300-48MP.

Supported Browsers

Juniper Routing Director supports the latest version of the following browsers:

- Google Chrome
- Mozilla Firefox
- Safari

New Features

IN THIS SECTION

- Device Life-Cycle Management | 10
- Observability | 11
- Trust and Compliance | 14
- Service Orchestration | 14
- Network Optimization | 15
- Planner | 16
- Active Assurance | 17
- Juniper Routing Director Installation | 18
- Beta Features | 18
- Deprecated Features | 19

This section describes the features available in Juniper® Routing Director Release 2.8.0.

Device Life-Cycle Management

Device life-cycle management (LCM) extends over the entire life cycle of a device. As part of device LCM, you install the device onsite, bring the device under management, monitor the device when it is in production, and finally decommission the device.

Juniper Routing Director Release 2.8.0 extends device LCM to the following platforms:

- **Device Support—**The Routing Director 2.8.0 release supports the following devices:
  - MX301
  - SRX4000 Series—SRX4100, SRX4200, SRX4300, SRX4600, and SRX4700.

- EX4400 Virtual Chassis

You can onboard devices, export device details, view audit logs, and upgrade software. See Supported Devices for a complete list of devices supported by Routing Director.

- **Adopt brownfield Junos devices using the management VRF**—Adopt brownfield Junos devices using the **Use Management VRF** toggle button on the Add a Device page (**Inventory > Devices > Network Inventory > Add Devices**). When you enable the button, the device uses the `mgmt_junos` VRF to connect to Routing Director to separate management traffic from data traffic.

  [See Add a Device to Routing Director.]

- **Support for gNMI dial-in connection for telemetry**—Routing Director uses gRPC Network Management Interface (gNMI) dial-in connectivity to obtain device telemetry if you onboard the device using a network implementation plan.

  Routing Director acts as the gNMI client and initiates the gNMI connection to the devices. Earlier, the devices acted as the gNMI clients and initiated the connection to Routing Director (dial-out connectivity).

  When you upgrade from an earlier release to Routing Director Release 2.8.0, all the dial-out connections change to dial-in connections without any data loss.

  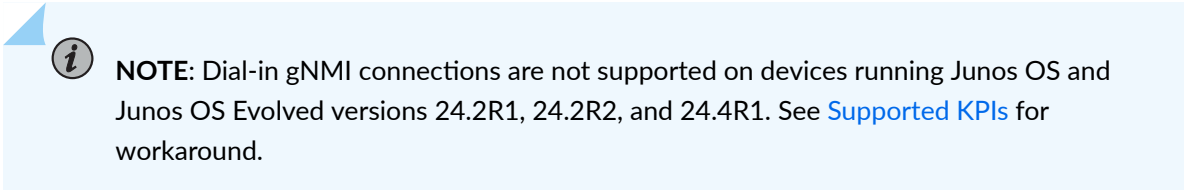

  NOTE: Dial-in gNMI connections are not supported on devices running Junos OS and Junos OS Evolved versions 24.2R1, 24.2R2, and 24.4R1. See Supported KPIs for workaround.

  [See Firewall Requirements.]

Observability

You can use Routing Director to view your entire network topology in real time and monitor network health. Additionally, you receive notifications about network anomalies and troubleshooting guidance.

With observability, Routing Director monitors and analyzes the network and its components by using key performance indicators (KPIs), device logs, and metrics. Observability includes alerts and alarms that notify you about network issues.

Routing Director also runs connectivity tests using synthetic traffic to identify connection issues between devices in your network. In addition, the real-time routing dashboard allows you to actively monitor the overall routing health of your network. Timely detection of anomalies enables you to take prompt action and minimize the impact of any issues.

Juniper Routing Director Release 2.8.0 provides the following additional observability feature:

- **AI/ML-driven device health monitoring on additional platforms**—In addition to ACX7024 and PTX10008, Routing Director monitors health and temperature for the following devices:
  - ACX7509
  - ACX7348
  - PTX10001-36MR
  - PTX10004
- **Detect physical layer faults in Routing Director using AI-ML**—Routing Director generates alerts when it detects physical layer faults during device onboarding or on an operational device. You can view the number of alerts on the Connectivity accordion of:
  - The *Device-Name* page (**Observability > Troubleshoot Devices > Device > Overview tab**) when a device is operational.
  - The *Device-Name* page (**Inventory > Device Onboarding > Onboarding Dashboard**) during device onboarding.

  You can view details of the alerts on the Relevant Events section of the accordion.

  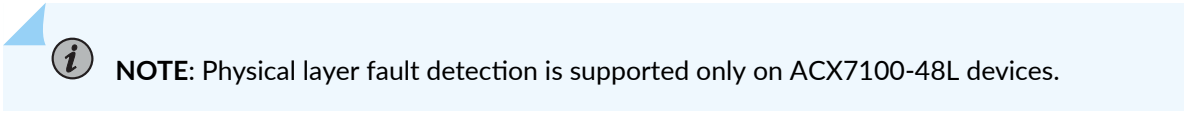

  NOTE: Physical layer fault detection is supported only on ACX7100-48L devices.

  [See Detect Physical Layer Faults.]
- **Detection of traffic loss and blackhole on MX Series devices using AI-ML**—Routing Director detects traffic loss and blackholes on MX204, MX240, and MX10008 devices in addition to the PTX10004 and PTX10008 Series devices. Routing Director displays this information on:
  - Routing and MPLS accordion of the Overview tab of the *Device -Name* page (**Observability > Troubleshoot Devices > *Device -Name***).
  - Alerts tab of the Events page (**Observability > Events**).

  [See Detect Blackholes.]
- **Increase KPI Data Retention**—You can now increase the KPI data retention period in Routing Director beyond the default one-week limit. After you configure the retention period, the date-time selector automatically adjusts. This helps you query KPI data for the newly configured range on the following pages:
  - *Component-Name* accordion (**Observability > Troubleshoot Devices > *Device-Name* > Hardware/ Interfaces/ Routing & MPLS accordion > *Component-Name-link***)

- *Component-Name* accordion (**Orchestration > Instances > *Service-Instance-Name* Details page > Passive Assurance > *Component-Name-link***)
- *Instance-Name* page (**Observability > Health > Custom KPI Collection > Rule Instantiations > *Instance-Name***)

You can now query KPI data within the updated retention range.

> ⚠ **CAUTION**: Increasing the retention period requires additional disk space. The exact storage requirements depend on the configured retention duration and the volume of KPI data collected. For an estimate of the additional disk space needed, contact your Juniper Sales Representative.

[See Retention Policies in Routing Director.]

- **Use Export Manager to stream data to external systems**—Use Routing Director to stream operational and observability data, including underlay and device KPIs and syslogs to an external Kafka system. You can configure Kafka destinations, define streaming parameters, map data categories to pre-created Kafka topics, and add, pause, or resume streams while preserving configuration state.

  Configure TLS and SASL to secure communication between Routing Director and external Kafka system. Routing Director can export events and KPIs in continuous streams, enabling integration with third-party systems such as SIEM platforms, AI/ML pipelines, and other event-monitoring systems. A new menu, **Management**, is added to the menu bar on the left-side of the Routing Director GUI. Use the Export Manager page (**Management > Export Manager**) to configure data streams and the Kafka destination.

  [See Export Manager Overview.]

- **IGP anomaly alert generation**—Routing Director detects and displays all IGP anomalies as alerts on the Events page (**Observability > Health > Events**). These anomaly alerts directly contribute to the overall routing health computation displayed on the Health Dashboard (**Observability > Health > Health Dashboard**).

  Use these insights to correlate topology-level issues with overall routing health status, which enables faster diagnosis and improved visibility into the impact of IGP anomalies on network performance.

  [See IGP Anomaly Detection Overview.]

- **View BGP Route Details**—View all routes and paths advertised by BGP in Routing Information Bases (RIBs) of devices from the **Routes Tab** (**Observability > Routing > Route Explorer > Route Status**). You can view information such as the list of devices, prefix addresses and lengths, extended communities, details of the RIB where a route is installed, and so on.

  [See About the Routing Status Tab.]

Trust and Compliance

Routing Director helps protect the network from threats and vulnerabilities by periodically checking whether a target's configuration, integrity, and performance comply with predefined security benchmarks. The term *target* refers to a device or a device component. Routing Director distills the outcomes of these checks into a single trust score that you can use to determine how trustworthy a device is.

There are no new features in Release 2.8.0.

Service Orchestration

Service orchestration is the process of designing, configuring, validating, deploying, and monitoring a network service. Routing Director automates the entire life cycle of a network service by providing workflows that execute the tasks required to deliver a service. You can provision various network services by using predefined service designs. The service catalog is an inventory of service designs, which are templates that provide guidelines and parameters for instantiating a service. A service instance defines the elements of a service. A service order includes the instruction to create, modify, or delete a service instance. After you initiate a service order and provision it, Routing Director activates the automated workflow to provision the service in the network. After provisioning, Routing Director automatically monitors network health and measures service quality.

Juniper Routing Director Release 2.8.0 provides the following additional service orchestration features:

- **Updated ESI configuration format in the L2 resource instance**—You can define the prefix and suffix for the Ethernet Segment Identifier (ESI) ID, and the start and end values for the ESI ID range on the Modify *L2-Addr-Resource-Instance-Name* page (**Orchestration > Service > Resource Instances > L2-Addr**). These values create a structured pool of ESI IDs that Routing Director uses when allocating placement ESI parameters for EVPN services.

  [See Configure Resource Pools for Resource Instances, Add EVPN Service Post Update Placements Parameters, and Add EVPN-VPWS Service Post Update Placements Parameters.]

- **Configure custom local VPWS ID**—Define a custom local VPWS ID for each CE device-facing site network access on the Add Connection page of the Add E-Line EVPN VPWS CSM wizard (**Orchestration > Service > Instances > E-Line EVPN VPWS CSM**). The local VPWS ID of one CE-device facing site network access is applied as the remote VPWS ID in the second CE device configuration. If no value is provided, the system automatically assigns default VPWS IDs 1 and 2 to the first and second CE device-facing site network accesses. This ensures consistent VPWS ID mapping across CE device-facing site network accesses.

  [See Add EVPN-VPWS Site and Site Network Access Details.]

- **Upgrade service design versions in bulk**—You can upgrade multiple service design versions in bulk from the Service Designs page (**Orchestration > Service Catalog > Service Designs**). After a successful

upgrade, the **Version** column displays the upgraded version of the service designs used by service instances.

[See About the Service Designs Page.]

Network Optimization

Routing Director's network optimization use case helps you optimize resource utilization, boost network performance, and ensure reliable, efficient data delivery. By using an intent-based approach, Routing Director optimizes the network through active life cycle management of label-switched paths (LSPs).

You can create a path intent using the Routing Director GUI. Path intents are specific LSP configurations that define how traffic is steered through the network. In traditional methods, you configure and provision each path in a tunnel individually with all its attributes. With path intent, you can create sub-profiles of attributes that can be reused for creating paths. This modular approach reduces redundancy and streamlines the process of provisioning multiple tunnels.

When you apply the path intent to the network, Routing Director interprets these intent-based sub-profiles and automates the creation, modification, and deletion of tunnels and LSPs. By autonomously executing the required actions, Routing Director aligns the network state with the specified intent. Routing Director ensures that LSPs are established based on network policies, traffic engineering constraints, and service level agreements (SLAs).

Juniper Routing Director Release 2.8.0 provides the following additional network optimization features:

- **Re-parse device collection**—To manually trigger re-parsing of previously-collected configuration and `show` commands, click **Re-parse** on the Advance tab of the Topology Settings page (**Observability > Network > Topology > Topology Menu Bar > Settings** icon).

  Use this option when the network model configuration state is out of sync with the collected configuration file and operational command output.

  [See View Network Topology Details.]

- **Map display options on the Topology page**—Use **Show World Map** on the Map tab of the Topology Settings page (**Observability > Network > Topology > Topology Menu Bar > Settings** icon) to select the map type on the Topology page. In case of air-gapped installation, you can choose a map option that uses local map data from Routing Director and does not require internet access.

  [See View Network Topology Details.]

- **Flexible metric selection for the color legend**—Select different traffic metrics on the Links tab of the Topology Settings page (**Observability > Network > Topology > Topology Menu Bar > Settings** icon). The selected metric appears in the color legend at the bottom right corner of the Topology map. This helps you to quickly analyze various utilization types.

  [See View Network Topology Details.]

- **Configure multiple autonomous systems**—Routing Director allows each BGP-LS peer address to have a different AS number. You can also configure multiple BGP-LS peers per AS. This granular control over AS assignments enhances network segmentation, isolates routing policies, and tailors routing strategies according to operational needs.

  Use the Dynamic Topology tab of the Topology Settings page (**Observability > Network > Topology > Topology Menu Bar > Settings** icon) to configure multiple BGP LS peers.

  [See View Network Topology Details.]

Planner

Planner is used for offline visualization and detailed architectural planning of any production network. Planner enables you to forecast the impact of changes to your network, such as additional traffic, shifts in traffic flows, and new capacity or services.

Planner generates a topology view of a network, enabling you to add, remove, and reconfigure network elements. Using the network topology view, you can model and visualize dynamic, explicit routing paths, designed to operate within end-user defined constraints. The effects of these changes and other traffic scenarios can be simulated without affecting the production network.

Juniper Routing Director Release 2.8.0 provides the following additional planner feature:

- **View offline topology changes**—View the changes you have made to your offline network (model). Run **Update Model and Paths** and click the **Reload** icon on the Topology page (**Planning > Networks > Offline Models > *Model-Name* > Open**).
  Click **Update Model and Paths** on the Topology page to save your changes and select whether Routing Director must:

  - Recompute all paths
  - Recompute newly-added paths, or
  - Only update parameters

  Use **Reload** to refresh the topology view.

  [See About the Offline Topology Page.]

- **Group working models using tags**—Use tags to organize and group working models. A tag is a label in the *key:value* format. Tags enable faster identification of related working models.
  You can add one or more tags to a working model on the Working Models page (**Planning > Networks > Offline Models**).

  [See About the Working Models Page.]

- **View reports in the Topology page**—View the What-If failure simulation report within the Topology page (**Planning > Networks > Offline Models > *Model-Name* > Open**).

  When you select a link, tunnel, or demand in the report, the topology map displays the original and rerouted paths, link utilization-related changes, and so on. This visualization helps you to quickly evaluate the impact of the simulated failure and plan mitigation.

  [See About the Offline Topology Page.]

Active Assurance

Active Assurance is a programmable test and monitoring solution, which generates synthetic traffic in the underlay network to gain continuous insights on network quality, availability, and performance. Active Assurance uses Test Agents, which are measurement points in your network. Test Agents generate and receive synthetic traffic, and enable you to continuously monitor and validate the infrastructure. You can deploy Test Agents at strategic locations in your network and install them on routers running Junos® OS Evolved, x86 hardware, or on virtual machines (VMs). Routing Director uses RPM to collect metrics data for Juniper Networks® MX Series Universal Routers and Juniper Networks® PTX Series Routers.

Juniper Routing Director Release 2.8.0 provides the following additional Active Assurance features:

- **Centralized view of Measurement reports**—View the reports generated by Measurements on the Reports Tab of the Measurement Details page **(Observability > Active Assurance > Measurement Explorer**). The centralized view helps you:

  - Verify outcomes by viewing Measurement validation results in one place.
  - Troubleshoot issues faster using plugin-specific configuration data.
  - Monitor the overall health, performance, and behavior of the Measurement using general details such as status, duration, start and end time, and so on.

  Reports are generated only after a Measurement reaches its final state (Passed, Failed, Error, or Stopped). Reports are available only for Measurements related to QoS Policy profiling and Path MTU discovery.

  NOTE: Reports are available only for Monitors. Tests do not generate reports.

  [See About the Reports Tab.]

- **Support for additional plug-ins**—Routing Director enables you to evaluate the QoS in your network using the following plug-ins:

  - **Cisco IP SLA Ping**—Use this plug-in to measure network responsiveness and verify IP (IPv4 and IPv6) reachability between Cisco devices. The plug-in sends ICMP Echo Requests to the

configured endpoint and evaluates metrics such as response time and packet loss. A verification test, using this plug in, helps detect issues related to connectivity, latency, and path stability.

- **Path MTU Discovery**—Use this plug-in to determine the maximum transmission unit (MTU) supported along a network path. The plug-in sends probe packets of varying sizes to identify the largest packet that can traverse the path without fragmentation. A verification test, using this plug in, helps detect MTU mismatches and fragmentation constraints that may impact performance.
- **QoS Policy Profiling**—Use this plug-in to validate Quality of Service (QoS) behavior and bandwidth distribution across all IPv4 and IPv6 traffic classes. This plug-in uses TCP and UDP flows to verify throughput allocation, delay characteristics, packet loss, and buffer behavior. A verification test, using this plug in, helps identify QoS misconfigurations, prioritization issues, and policy enforcement gaps.

[See Supported Plug-ins.]

Juniper Routing Director Installation

Juniper Routing Director Release 2.8.0 provides the following installation-related feature:

- **Install on a single node**—Install Routing Director on a single virtual machine (VM). The single VM functions as both a primary node and worker node.

  Install Routing Director on a single node in lab environments, POCs, and small scale deployments where only a minimal number of devices (under 10) need to be managed. The single-node deployment must be used only when scalability or availability requirements are far less stringent than those of production deployments.

  The installation process is the same as that of a multinode cluster. You can back up and restore a single-node setup, but you cannot upgrade an older release of Routing Director to a single-node setup.

  [See Routing Director Implementation, System Requirements, and Deploy the Cluster.]

Beta Features

Juniper Routing Director Release 2.8.0 provides Beta support for the following features:

- **Use Model Context Protocol (MCP) server to query Routing Director**—A network operator can use any AI agent such as Claude, Copilot, and ChatGPT to query Routing Director through an MCP server. The MCP server helps network operators to query network data, create dashboards, and access KPIs conversationally instead of writing code or learning complex API syntax.

  [See Query Routing Director Using MCP.]

- **Use LLM Connector to query Juniper documentation**—Use LLM Connector to query only the documentation available on the Documentation site (https://www.juniper.net/documentation/), and cite references for the responses. Configure LLM Connector to query documentation on the Documentation tab of the Configure LLM Connector widget (Settings Menu > System Settings > Organization Settings). in the Organization Settings (**Settings Menu > System Settings > LLM Connector**).

  [See LLM Connector Overview.]

- **Configure an IRB interface for L3VPN services**—You can configure an IRB interface for L3VPN services for the following scenarios:

  - L3VPN with EVPN using regular untagged interfaces with OSPF as PE-CE protocol and insights.
  - L3VPN with EVPN using regular untagged interfaces with BGP as PE-CE protocol and insights.
  - L3VPN with EVPN using interfaces in VLAN mode with OSPF as PE-CE protocol and insights.
  - L3VPN with EVPN using interfaces in VLAN mode with BGP as PE-CE protocol and insights.

  [See Add L3VPN Site and Site Network Access Details and Add L3VPN Service Post Update Placements Site Network Access Parameters.]

- **Upload a customized service design**—Upload customized service designs to Routing Director by using the service orchestration cMGD CLI.

  > **NOTE**: To create and customize service designs, contact Juniper Networks Professional Services.

  You can view the uploaded service designs on the Service Designs page (**Orchestration > Service > Service Catalog**) and use them to provision corresponding services in the network.

  [See Upload a Customized Service Design.]

Deprecated Features

The following features are deprecated in Juniper Routing Director Release 2.8.0.

- **upgrade_routing-director-release-build-ID.tgz compressed archive file**—You can no longer upgrade to the latest release of Routing Director using the **.tgz** format file. You must use the **.img** file to upgrade to the latest release.

  [See Upgrade Routing Director.]

- **Manual creation of topology resource instances**—You cannot create topology resource pools from the Resource Instances page (**Orchestration > Service > Resource Instances > + Add New Resource**

Instance). You can modify only existing topology resource pools from the Resource Instances page. Configure topology resources by specifying them in the network implementation plan.

[See Configure Resource Pools for Resource Instances.]

- **EVPN ESI name and count configuration**—The legacy EVPN ESI section in which you configured a name and count for EVPN ESIs is deprecated. This section is available only to support service instances provisioned in releases prior to release 2.8.0. We recommend that you migrate service instances from earlier releases to the new EVPN ESI configuration in which you must enter a fixed prefix, suffix and a variable middle range.

  [See Configure Resource Pools for Resource Instances.]

Known Issues

IN THIS SECTION

- Device Life-Cycle Management | **20**
- Observability | **21**
- Service Orchestration | **29**
- Active Assurance | **31**
- Network Optimization | **34**
- Network Planner | **37**
- Trust | **38**
- Administration | **38**
- Installation and Upgrade | **38**

This section lists the known issues in Juniper Routing Director.

Device Life-Cycle Management

- Changing the router ID of a device after onboarding might create a duplicate node in the topology.

  Workaround: If you want to change the router ID, you must offboard the device, update the router ID in the configuration, and then onboard the device again.

- When a vMX is deployed using I2C ID 161, all commit operations will fail after a subscriber is created.

  Workaround: Delete the subscribers and then commit the configuration.

- Routing Director triggers the configuration templates included in a device profile and interface profile only during the initial onboarding of the device. You cannot use the configuration templates included in the device profiles and interface profiles to apply additional configuration on a device after the device is onboarded.

  Workaround: If you need to apply additional configuration on a device after the device is onboarded, you need to manually apply the configuration using the CLI or by executing the configuration templates through the Routing Director GUI.

Observability

- Inconsistent Y-axis scaling in the **Input, Output and Drop Rate** graph on the Traffic Loss page (**Observability > Health > Troubleshoot Devices > Device-Name > Overview > Routing and MPLS accordion > Traffic Loss Alert** link) causes a misleading representation of traffic rates.

  Workaround: None.

- When an adjacency fails or is established, Routing Observability detects these events and reports them on the Events page (**Observability > Health > Events**). However, the same adjacency event may be reported multiple times.

  Workaround: Treat these repeated events as a single event by correlating the self and neighbor ID combination.

- When both L1 and L2 are enabled simultaneously, the IGP Prefixes tab (**Observability > Routing > Route Topology**) does not display prefixes for both levels. Instead, only L1 prefixes are displayed. Dual-level (L1 + L2) prefix support is not supported.

  Workaround: To view prefixes correctly, on the Analytics tab of the Device Profile page (**Inventory > Devices > Device and Interface Profiles > Add > Device Profile**), select **LSDB** for a device in L1 or in L2 only. Avoid selecting a device configured for both L1 and L2.

- When devices under monitoring fail catastrophically and are subsequently offboarded from Routing Directory, these stale devices may still show up in the Route Explorer's devices table (**Observability > Routing > Devices tab**).

  Workaround: Remove stale devices by manually deleting them from the Routing Director's internal database.

- After restoring a backup, the Adjacencies tab (**Observability > Routing > Route Explorer**) may display an incorrect state for BGP peers.

Workaround: Flap the BGP peers on the monitored device.

- Under certain conditions, transient traffic loss may cause a Packet Drop (Major) alert to be misclassified as a Blackhole (Critical) alert. You can identify such alerts by checking for identical start and end times.

  Workaround: None.

- Alerts are not displayed on the Relevant Events section of all accordions on the Passive Assurance tab (**Orchestration > Instances > Service Instances > *Service-Instance-Name* hyperlink**).

  Workaround: You can view the alerts on the Events page (**Observability > Event**) or on the respective graphs.

- Although the **Documentation** mode is configured, the **Conversation** mode is erroneously activated when you open an LLM Connector chat window.

  Workaround: Switch the interaction mode to *Documentation* manually before initiating a prompt. Once selected, the session operates as expected.

- Even with auto-refresh enabled, the alerts listed and the data shown in the graph on the IS-IS Routing Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS > IS-IS > IS-IS Adjacency Flap**) may not be synchronized.

  Workaround: Close and reopen the IS-IS Routing Details for *Device-Name* page to see the latest data.

- **SASL Username** and **Password** fields are cleared unexpectedly when you enable the **Use TLS Encryption** toggle button on the Add Destination page (**Management > Export Manager > Manage Destination**).

  Workaround: Re-enter the **SASL Username** and **Password** fields after you enable the **Use TLS Encryption** toggle.

- In spite of auto-refresh, alerts listed and data displayed on the graph might not be synchronized on the Output Traffic Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Output Traffic**).

  Workaround: Close and reopen the Output Traffic Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Output Traffic**) to see the latest data.

- Starting with Release 2.8.0, Juniper Routing Director initiates the gNMI dial-in connection to Juniper devices to collect telemetry data. You must ensure that the corresponding firewall rules allow traffic towards devices on port 32767 from the Routing Director collector.

  We have disabled certificate verification by default. If you want to re-enable certificate validation, ensure that the devices in your network run on any releases other than Junos OS or Junos OS

Evolved Releases 24.2R1, 24.2R2, and 24.4R1. Use the following REST API command to re-enable certificate validation:

```
curl -X PUT -H "Content-Type: application/json" \
    -u test@test.com:Test-Password \
    "https://VIP-1/api/v1/orgs/{org-id}/gnmi/options" \
    --data '{"device": {"client-certificate-request": "require-certificate-and-verify"}}' --
insecure
```

- Sometimes, there is a mismatch in the severity level that is displayed in the Severity column of the Troubleshoot Devices page (**Observability > Health > Troubleshoot Devices**) and the Device tab of the Topology page (**Observability > Network > Topology**).

  Workaround: None.

- The Traffic Loss link on the Routing and MPLS accordion (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview** tab) shows zero alerts when there is a delay in raising an alert due to slow data processing. This may result in an inaccurate active alert count.

  The issue gets resolved when the data is processed. You can see a non-zero active alert count again when the data is processed.

  Workaround: None

- You might see duplicate blackhole alerts on the Traffic Loss page (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Routing and MPLS** accordion > click **Traffic Loss Alert** link) even though traffic is continuous. This issue occurs only when you upgrade from Release 2.5.0 to Release 2.6.0 or Release 2.7.0. You won't encounter this issue if you have installed Release 2.7.0 afresh.

  Workaround: If you have upgraded to release 2.7.0, delete the blackhole index, by running the `curl -X DELETE "http://$(kubectl get svc -n common opensearch-cluster-master -ojsonpath='{.spec.clusterIP}'):9200/blackhole_detection` command.

  NOTE: You might lose all the data processed before the upgrade.

- When KPIs continuously oscillate between fixed values in a repeating pattern, the boundary initially adapts as expected, but after a few hours, it begins to readjust even though the oscillation pattern remains unchanged. This behavior can persist even after full adaptation, causing the boundary to continue oscillating unnecessarily.

  Workaround: None.

- In setups with parallel links between two nodes, adjacency or link flaps are reported only on a complete loss or restoration of connectivity between the nodes, rather than on individual link transitions.

  Workaround: None

- There is an unexpected delay in reporting an anomaly related to a sudden decrease in nodes. The anomaly is reflected only after the total delay period has passed.

  Workaround: You can use the REST API to view information related to this anomaly.

- The Device Count column on the IGP Prefixes tab (**Observability > Routing > Route Topology**) might display inaccurate data. It takes approximately 30 minutes for the device count to be updated when a new device starts originating the prefix or when a device stops originating the prefix.

  Workaround. Device count may not immediately reflect recent changes and may be inaccurate for up to 30 minutes. We recommend that you wait for approximately 30 minutes to get the latest device count.

- When the responses are delayed, the LLM Connector chat window auto-scrolls up and down unexpectedly.

  Workaround: None.

- After a device is onboarded, Routing Director continuously monitors the KPIs related to device health. For each KPI, Routing Director monitors the KPI, forecasts the range, and detects any anomalies that occur. If a KPI value changes, the forecasted range takes approximately two hours to stabilize.

  Workaround: None.

- While adding a device profile for a network implementation plan, if you enable Routing Protocol Analytics then the routing data is collected for the devices listed in the device profile. When you publish the network implementation plan, even though the onboarding workflow appears to be successful there might be errors related to the collection of routing data for these devices. Because of these errors, the devices will not be configured to send data to Routing Director and therefore the routing data will not be displayed on Route Explorer page of the Routing Director GUI. This issue occurs while offboarding devices as well, where the offboarded devices continue to send data to Routing Director.

  This issue also occurs when you have not configured ASN or Router ID on the devices, or when you have locked device configuration for exclusive editing.

  Workaround: To fix this issue:

  1. Do one of the following:

- Check the service logs by running the `request paragon debug logs namespace routingbot app routingbot service routingbot-apiserver` Shell command. Take the necessary action based on the error messages that you see in Table 6 on page 25.

Table 6: Error Messages

Error Messages	Issue
Failed to get device profile info for dev_id {dev_id}: {res.status_code} - {res.text} Failed to get device info for dev_id {dev['dev_id']}. Skipping device.	The API call to PAPI to get the device information has failed.
No results found in the response for dev_id {dev_id} Failed to get device info for dev_id {dev['dev_id']}. Skipping device.	The API call to PAPI returns a response with no data.
Complete device info not found in the response for dev_id {dev_id} : {device_info}	The API call to PAPI returns a response with incomplete data.
No data found for dev_id {dev_id} from PF	The API call to Pathfinder to get the device information has failed.
Required data not found for dev_id {dev_id} from PF data:{node_data}	The API call to Pathfinder to get device information returns a response with incomplete data.
EMS config failed with error, for config: {cfg_data} or EMS Config push error {res} {res.text} | try: {retries}. Failed to configure BMP on device {mac_id}	BGP configuration has failed.
Invalid format for major, minor, or release version : {os_version}	The device's OS version is not supported.
Error POST {self.config_server_path}/api/v2/config/device/{dev_id}/ {data} {res.json()}	Playbook application has failed.

Table 6: Error Messages *(Continued)*

Error Messages	Issue
Error PUT:{self.config_server_path}/api/v2/config/device/{dev_id}/ {data} {res_put.json()}	Playbook removal has failed.
Error PUT:{self.config_server_path}/api/v2/config/device/{dev_id}/ {data} {res_put.json()}	Device or playbook application to device-group has failed.
Error PUT {self.config_server_path}/api/v2/config/device-group/{site_id}/ {data} {res_put.json()}	Device or playbook removal from device-group has failed.

- Examine the device configuration to check whether the device shows unexpected absence or presence of the configuration. For example, you can,
  - View the configurations present under `set groups paragon-routing-bgp-analytics routing-options bmp`.
  - Check the device configuration in the JTIMON pod.

2. After resolving the above issues, edit the device profile of the network implementation plan that you have applied for the device. Based on whether you are onboarding or offboarding devices, enable or disable the Routing Protocol Analytics option in the device profile.
3. Publish the network implementation plan.
4. Verify whether the required results are seen based on the data that is displayed on the Route Explorer page of the Routing Director GUI.

- On the Interfaces accordion, FEC uncorrected errors charts are available only on interfaces that support speeds equal to or greater than 100-Gbps.
- After you apply a new configuration for a device, the Active Configuration for *Device-Name* page (**Observability> Troubleshoot Device > *Device-Name* > Configuration accordion > View active config link**) does not display the latest configuration immediately. It takes several minutes for the latest changes to be reflected on the Active Configuration for *Device-Name* page.

  Workaround: You can verify whether the new configurations are applied to the device by logging in to the device using CLI.
- The number of unhealthy devices listed on the Troubleshoot Devices and Health Dashboard pages (**Observability > Health**) do not match.

Workaround: None.

- You cannot delete unwanted nodes and links from the Routing Director GUI.

  Workaround: Use the following REST APIs to delete nodes and links:

  - REST API to delete a link:

    [DELETE] https://{{server_ip}}/topology/api/v1/orgs/{{org_id}}/{{topo_id}}/links/{{link_id}}

    

    NOTE: You can follow the steps described "here" on page 28 to get the actual URL.

    For example,

    - URL: 'https://10.56.3.16/topology/api/v1/orgs/f9e9235b-37f1-43e7-9153-e88350ed1e15/10/links/15'
    - Curl:

```
curl --location --request DELETE 'https://10.56.3.16:443/topology/api/v1/orgs/
f9e9235b-37f1-43e7-9153-e88350ed1e15/10/links/15' \
--header 'Content-Type: application/json' \
--header 'Authorization: Basic dGVzdDFAdGVzdC5jb206RW1iZTFtcGxz'
```

  - REST API to delete a node:

    [DELETE] https:// {{*Server_IP*}}/topology/api/v1/orgs/{{*Org_ID*}}/{{*Topo_ID*}}/nodes/{{*Node_ID*}}

    

    NOTE: You can follow the steps described "here" on page 28 to get the actual URL.

    For examples,

    - URL: ' https://10.56.3.16/topology/api/v1/orgs/f9e9235b-37f1-43e7-9153-e88350ed1e15/10/nodes/1'
    - Curl:

```
curl --location --request DELETE 'https://10.56.3.16:443/topology/api/v1/orgs/
f9e9235b-37f1-43e7-9153-e88350ed1e15/10/nodes/11' \
--header 'Content-Type: application/json' \
--header 'Authorization: Basic dGVzdDFAdGVzdC5jb206RW1iZTFtcGxz' \
```

Use the following procedure to get the actual URL that you use in CURL for deleting a link or a node:

1. Navigate to the Topology page (**Observability > Topology**).
2. Open the developer tool in the browser by using the **CTRL + Shift + I** buttons in the keyboard.
3. In the developers tool, select **Network** and select the **XHR** filter option.
4. Identify the link index number or node number. To identify the link index number to the node number:
   a. On the Topology page of the Routing Director GUI, double click the link or the node that you want to delete.

      The Link *Link-Name* page or the Node *Node-Name* page appears.

   b. Navigate to the Details tab and note the link index number or the node number that is displayed.
5. In the developers tool, select and click the row based on the link index number or the node number that is related to the link or the node that you want to delete.
6. Copy the URL that you need to use to delete the link or node in CURL.

- Not all optics modules support all the optics-related KPIs. See Table 7 on page 28 for more information.

  Workaround: None.

  Table 7: KPIs Supported for Optics Modules

Module	Rx Loss of Signal KPI	Tx Loss of Signal KPI	Laser Disabled KPI
SFP optics	No	No	No
CFP optics	Yes	No	No
CFP_LH_ACO optics	Yes	No	No
QSFP optics	Yes	Yes	Yes
CXP optics	Yes	Yes	No

Table 7: KPIs Supported for Optics Modules *(Continued)*

Module	Rx Loss of Signal KPI	Tx Loss of Signal KPI	Laser Disabled KPI
XFP optics	No	No	No

- For PTX100002 devices, the following issues are observed on the Interface accordion (**Observability > Health > Troubleshoot Devices > Device-Name > Overview**):
  - On the Pluggables Details for *Device-Name* page (**Interfaces accordion > Pluggables data-link**), the Optical Tx Power and Optical Rx Power graphs do not display any data.
  - On the Input Traffic Details for *Device-Name* page (**Interfaces accordion > Input Traffic data-link**), the Signal Functionality graph does not display any data.

Service Orchestration

- MX104 and ACX2200 devices do not support gNMI rules. Therefore, the Passive Assurance tab (**Orchestration > Instances service-instance-name hyperlink > Service-Instance-Name Details**) does not display Physical Interfaces and Logical Interfaces accordions for these devices.

  Workaround: None.

- If you use service designs created in Release 2.6.0 or earlier, we recommend that you upgrade your service designs to the latest version.

  Service configurations previously hard-coded in the GUI for older service designs are now automatically generated based on the service design definition. If you continue to use the older service design, ensure that you manually enter the IPv4 loopback address using the GUI.

- If you try to delete (deprovision) a service instance that previously ran in Dry-Run mode, the system mistakenly executes the delete action in Dry-Run mode. As a result, the service instance is not removed, and its status remains unchanged.

  Workaround: Do one of the following:

  - Run **Provision** on the instance to reset its state. Once provisioning is complete, run **Deprovision**.
  - Use the `DELETE /service-orchestration/api/v1/orgs/{orgId}/order/customers/{custId}/instances/{instId}` REST API to bypass the Dry-Run behavior.

- Not all devices are listed in the L3VPN accordion on the Passive Assurance tab (**Orchestration > Instances > *service-instance-name* hyperlink > Service-Instance-Name Details**).

Workaround: The issue primarily occurs during upgrades. Restart all *iTSDB* pods simultaneously to resolve the issue.

- When you configure an EVPN service instance, both **mpls-evpn** and **pbb-evpn** VPN service types are displayed as options. However, only **mpls-evpn** is supported on Routing Director GUI and is the default service type.

  Workaround: None.

- On the Resource Instances page (**Orchestration > Service > Resource Instances**), the **network-operator:topo** resource is a system-managed resource. As a result, the Workflow Run ID column may be empty when the system generates the resources. The Workflow Run ID column is set only if you click the **Update** button.

  Workaround: None.

- The description for the interfaces is missing in the **Description** column of the Add or Edit Devices section of the Add Network Implementation Plan page (**Inventory > Device Onboarding > Network Implementation Plan > +**).

  Note that the description for sub-units will be the same as that of the main interface description. If the descriptions for sub-units et-0/0/9.100 and et-0/0/9.200 are missing, you can refer to the description of the main interface, et-0/0/9.

  Workaround: None.

- In rare high-load scenarios, provisioning of an EVPN instance may fail due to the unavailability of temporary back-end resources.

  Workaround: Try provisioning the service again.

- The following accordions on the Passive Assurance tab (**Orchestration > Instances > *Service-Order-Name* Details**) displays incorrect or no data:

  - BGP accordion—The VPN State column displays incorrect data for customer edge (CE) or provider edge (PE) devices with IPv4 or IPv6 neighbors.

  - OSPF accordion—There are no IPv6 entries in the Neighbor Address column for CE or PE devices with IPv6 neighbors.

  - L3VPN accordion—The VPN State column displays incorrect data for OSPF and BGP protocols. The Neighbor Session and VPN State columns are blank for CE or PE devices with static IPv4 or IPv6 address.

  This issue occurs only for an L3VPN service.

  Workaround: None.

- For an MX 240 device, the OSPF-related data is not populated on the Passive Assurance tab (**Orchestration > Instances > *Service-Order-Name* Details**).

  Workaround: Configure OSPF on the customer edge (CE) device.

- When you click the **Refresh** icon on the *Service-Instance-Name* Details page (**Orchestration > Instances > *Service-Instance-Name***), you may not see the latest events in the Relevant Events section.

  Workaround: To view the latest events, instead of using the Refresh icon go to the Service Instance page (**Orchestration > Instances**) and select the service instance for which you need to see the latest events.

- The Order History tab on the *L3VPN-Name* Details page (**Orchestration > Instances > *Service-Instance-Name*** hyperlink) lists all the order history if you deprovision a service instance and later provision a service using the same details as that of the deprovisioned service.

  Workaround: None.

- In a scaled setup, you cannot upgrade service instances in bulk.

  Workaround: We recommend that you upgrade only a few service instances (less than 4) at a time.

Active Assurance

- If you upload an invalid plugin in the Plugin Inventory page (**Inventory > Active Assurance**), the UI does not display a clear error message and remains stuck at 0%.

  Workaround: Upload a valid **.nap** file.

- When you download a report for a Monitor that has a large number of Tasks (approximately 100), the report contains incomplete data, and the Measurement Summary section in the report omits some tasks. The report does not render the ten most recent events or the event bar for each measurement.

  Workaround: None.

- The results produced by the Test Agent are impacted if the Test Agent clock has a large offset. That is, if the local time is in the past or future. This means:

  - Timestamps for Metrics for any Stream produced by a Measurement running on that Test Agent are affected.
  - Event activation time and event deactivation time are affected.

  Therefore, it can result in the incorrect evaluation of Test Execution, as Metrics or Events are not included in the time range the system considers as the Test Execution run time. You may not be explicitly warned about this situation. However, the issue manifests with time shifted metrics or events.

Workaround: Time synchronization is a requirement for Test Agents. Ensure Test Agent clock is synchronized.

- When you run a QoS profiling test, the TCP throughput is affected by congestion control.

  QoS policy profiling reflects actual network behavior. So, TCP throughput observed during QoS profiling tests is influenced by TCP congestion control. Drop policers can lower measured TCP performance because packet loss triggers congestion responses and reduces the sending rate. If you see reduced throughput in profiling results, we recommend reviewing the policers in use and consider using a traffic shaper instead. Shapers queue excess packets rather than dropping them, allowing profiling tests to represent the network's true capacity and performance more accurately.

- Only one path maximum transmission unit (Path MTU) server–client measurement pair can run at a time on the same server agent interface. When multiple Path MTU server measurements are started on the same agent interface, only the first server starts successfully. Subsequent servers produce error reports with the message: `Failed to create test socket: Address already in use`.

  Currently, when you create a Task for a Test execution, the GUI allows you to select one Server Agent and multiple client agents for the Path MTU plug-in, and therefore, you might encounter this issue.

  Workaround: None.

- If you reboot an ACX device that has a Test Agent installed, you might notice that Docker is removed and the Test Agent goes offline, affecting active assurance measurements.

  Workaround: Do the following:

  1. Log in to the router.
  2. Deactivate the *paa test-agent* service

     and commit the changes.

```
edit
root@paa-acx7100-1# deactivate services paa

[edit]
root@paa-acx7100-1# commit
commit complete
```

3. Reactivate the *paa test-agent* service and commit the changes.

```
[edit]
root@paa-acx7100-1# activate services paa

[edit]
root@paa-acx7100-1# commit
```

- In some rare cases, only Test Agents that are in an offline state and due for a plug-in upgrade are upgraded. The plug-in upgrade may not happen for Test Agents that are online and due for a plug-in upgrade.

  Workaround: Changing the active version of one of the plug-ins in the system (not necessarily the same plug-in or in the same organization) will make any pending upgrades to be revisited, causing the upgrade to continue for any online Test Agent pending to be upgraded. You can do this in one of the following ways:

  - You can use the Plugin Inventory page (**Inventory > Active Assurance**) to change the Active version of a plug-in back and forth between two plug-in versions.
  - Or, alternatively, use the API to re-enable the same plug-in version.

    1. Copy the `ID` of the plug-in version from the Plugin Inventory page.
    2. Run the following request to re-enable the same plug-in:

```
curl -H "Authorization: token $TOKEN" \
-X PATCH ${CCHOST}/active-assurance/api/v2/orgs/$ORG/plugins/${PLUGIN_ID}?
update_mask="enabled" \
-d '{
    "id": "'"${PLUGIN_ID}"'",
    "enabled": true
}' |jq
```

- The Metrics graph shows No Data for a Test that includes a Step with Measurements if:

  - The Test uses self-governed plugin.
  - If you click a Stream that produces Metrics while the Test is executing.

  This issue occurs if you set the same start time and end time.

Workaround: Ensure that you manually set the Custom Time Range to something meaningful instead. Once the Test execution is complete, the Metrics are shown correctly.

- When you restore a Routing Director instance, you might notice that some data such as Active Assurance Plugins and Packet Capture files may not be backed up. This is because backups are not done on any Kubernetes volumes.

  Workaround: We recommend that you download Packet capture files before you restore an instance and store them locally to analyze them. In case of Active Assurance Plug-ins, we recommend that you use the Plugin Inventory page (**Inventory > Active Assurance**) to upload the latest Plugin again on the new (restored) instance.

- After you take a backup and restore a Routing Director instance, some Test Agents might incorrectly display status as *Online*.

  Workaround: After the system is fully restored, perform the following steps:

  1. Restart `test-agent-gateway` one additional time to refresh Test Agents.
  2. Execute the `kubectl -n paa delete pod -l app=paa-test-agent-gateway` command from the Linux root shell.

- When you create a Monitor with 600 streams, you might encounter Monitor Creation Timeout error and the Monitor might automatically stop.

  Workaround: Restart the Monitor from the *Monitor-Name* page (**Observability > Active Assurance > Monitors > *Monitor-Name*) and click More > Start**) on the Routing Director GUI.

- The status of a Test Agent is shown as offline after the device's Routing Engine switches over from the primary Routing Engine to the backup Routing Engine, or vice versa. This issue occurs only if you are using a Junos OS version that is older than 23.4R2.

  Workaround: Reinstall Test Agent after the Routing Engine switchover.

- When you add a new host to the existing Monitor, the new measurements are not reflected in the Active Assurance tab of the Health Dashboard (**Observability > Health**).

  Workaround: None.

Network Optimization

- The Container LSP feature is not supported in Release 2.8.0. The **Container LSP Normalization** check box on the Organization settings page (**Settings Menu > System Settings**) and container LSP-related REST APIs in the API Reference Guide are non-functional.

  Workaround: None.

- The path computation engine (PCE) of Routing Director does not use the delay type chosen by users in the **Interface Delay Type** field (**Settings Menu > System Settings > Organization Settings > Network Optimization Settings**) for computing paths for tunnels that are delay-based or have a maximum delay constraint. Instead, the PCE always uses the average delay value.

  Workaround: None.

- Bulk deletion for Path Computation Element Protocol (PCEP) segment routing (SR) LSPs on Cisco IOS XR is not supported.

  Workaround: Delete the LSPs individually.

- When traffic on a label-switched path (LSP) drops to 0, Junos OS telemetry doesn't report a *0* value because `zero-suppression` is enabled by default. This can lead to unexpected results for features, such as bandwidth sizing and LSP rerouting on threshold crossing. The Routing Director's path computation engine continues to use the last known traffic value of the LSP during events requiring LSP rerouting.

  Workaround: Disable `zero-suppression` on the router by using the `set services analytics zero-suppression no-zero-suppression` command.

- Events listed on the Events page O**bservability > Network > Topology > Tunnels tab > Tunnel-Name > View > Event History)** do not contain route information of SR and SRv6 LSPs. Due to this issue, the **Show Path Changes** option might not work for SR and SRv6 LSPs.

  Workaround: None.

- When you remove the LSP delegation, the routing method is automatically changed from *default* to *routeByDevice*.

  Workaround: You must manually update the LSP routing method to the desired option.

- For bandwidth-sizing-enabled SR LSPs, the bandwidth resizing that is based on aggregate traffic through the LSP might not always happen as per the configured thresholds.

  There could be instances when an SR LSP's bandwidth gets changed, despite the aggregate LSP traffic value not exceeding the current bandwidth by the adjustment threshold percentage. There could also be instances when the LSP's bandwidth does not get resized, though the computed aggregate traffic differs from the current bandwidth by the configured threshold.

  This occurs due to incorrect comparison during bandwidth sizing of LSP traffic in accordance with the LSP bandwidth that is set while creating or modifying the LSP using Routing Director GUI or REST API.

- Incorrect RSVP link utilization can occur in the following scenarios:

  - Due to an LSP constraint, Routing Director reroutes a lower-traffic LSP instead of a higher-traffic LSP during Threshold Crossing Rerouting.

  - In the next path optimization, Routing Director removes the higher-traffic LSP's current path, including its bandwidth accounting along the path, which results in incorrect RSVP link utilization.

- Junos OS Release 22.4R1 and later have a limitation with SR-TE LSPs. For PCEP sessions to be established, you must disable the multipath feature using the following command:

  ```
  set protocols pcep disable-multipath-capability
  ```

  Secondary path is not supported.

- The status of SR-TE LSPs might be displayed as down after delegating or provisioning. There might be errors (RPD_SPRING_TE_ROUTE_LSP_MISMATCH) in RPD logs if there are parallel SR-TE LSPs (same source or destination nodes) using both node and adjacency SIDs.

  Workaround: All parallel SR-TE LSPs should use either node or adjacency SIDs.

- If you try to create an LSP using the REST API and if you are reusing an existing LSP name, then the REST API server does not return an error.

  Workaround: None.

- Not all columns in the Event History table (**Observability > Network > Topology > Tunnels tab > *Tunnel-Name* > View > Event History**) are applicable to event history. So, you might see blank columns.
  Workaround: None.

- After you perform a backup or restore operation, the traffic is displayed as 0 percent on the Topology page.

  Workaround: After the backup or restore operation, either restart the `pf-telemetry` pod or trigger a device collection., and also restart the `pcs` pod in `pf- namespace`.

- When a router which resides on an MPLS LSP path raises the overload (OL) bit in its IGP domain, sometimes, the tunnel may not get rerouted away from that router.

  Workaround: Manually specify a large delay on a node with the `overload` bit.

- Due to Kafka message size limitation, you can delete only 200 LSPs at a time.

  Workaround: None.

- If there are multiple ECMP diverse paths and if you have enabled periodic re-optimization, then the diverse LSPs might switch back and forth between two routing paths.

  Workaround: If you do not prefer this behavior, set the **Path Type** as Preferred on the Modify LSP page.

- Sometimes, an LSP provisioning might not be successful, and you might see the *PCC_Pending* error displayed on the tunnels table of the Topology **(Observability > Topology**) page.

Workaround: Restart the PCEP session on head-end routers by deactivating and activating the protocols and PCE-related statements in the Junos OS configuration.

- In broadcast links exist in the network, Segment Routing (SR) LSPs may not be created.

  Workaround: Change broadcast links to point-to-point links in the router configuration.

Network Planner

- If you modify the interface address and bandwidth from the Interfaces tab of the offline Topology page (**Planning > Networks > *Offline-Model* > Open**), then the changes are not reflected on the Links tab of the offline Topology page (**Planning > Networks > *Offline-Model* > Open**).

  Workaround: None.

- During path computation in Planner, links marked as down are still considered. Consequently, the computed path may include a down link.

  Workaround: None.

- Planner incorrectly uses TE metrics to compute SR tunnel paths when the routing method is set to *routebydevice*.

  Workaround: You can change the tunnel's routing method to ISIS or OSPF so that Routing Director can compute the tunnel path based on the IGP metric.

- On the Offline Topology page (**Planning > Networks > Offline Models > *Model-Name* > Open**), deleting a device does not automatically delete its associated links.

  Workaround: Manually delete links attached to the node before you delete a node.

- SR-LSP-related modeling requires you to create a network model using the **Import from Live** option. Currently, a network model created using **Import from Config** won't contain complete SR-related information.

  Workaround: Use **Import from Live** instead of **Import from Config** for SR networks.

- In an offline model, if there are links without interfaces, then the exhaustive failure simulation might fail and a report is not generated.

  Workaround: None.

- If you log out of the Routing Director GUI and re-login using the same tab or window of a browser, you may not be reconnected to Planning Offline Model or simulation notification services.

  Workaround: Refresh the browser before you re-login.

- The admin-group constraint that is set in a tunnel is not considered during the what-if failure simulation.

  Workaround: None.

- A tunnel and demand should not have the same name.. Otherwise, the status of the tunnel and demand might be displayed as *down*.

- In an offline model, only a primary LSP can be created. You cannot create a secondary LSP or a standby LSP in an existing offline model. You can, however, view the secondary or standby LSP-related details when you import a live network.

  Workaround: None.

Trust

There are no known issues in this release.

Administration

- In the case of a scale deployment, when the system is under resource stress, you may notice that the **papi-mon service** pod restarts. This happens as part of the self-healing recovery, so that the service is restored.

  Workaround: None.

- You cannot use a Transport Layer Security (TLS) certificate to onboard Nokia devices.

  Workaround: None.

Installation and Upgrade

- You might encounter the following error while upgrading or redeploying Routing Director configured with multiple network interface cards (NICs):

  ```
  Current host IP: 10.123.42.1
  Master Node 1 IP: 192.168.69.5
  error: This host is not master node 1, where this Deployment cluster was initially installed
  from!
  Please run upgrade from master node 1: 192.168.69.5
  ```

  This issue occurs because the **/root/epic/host.ip** file is populated with IP address used by the other NIC in the VM or is empty.

Workaround: Verify and re-populate **/root/epic/host.ip** with the appropriate IP address before upgrading or redeploying Routing Director.

You can also optionally make the **/root/epic/host.ip** file immutable to prevent overwrites by using `chattr +i /root/epic/host.ip`.

- Sometimes, it takes longer (approximately 10 minutes) to log in to the deployment shell of the cluster nodes. You see the following message before you can log in:

```
Awaiting configuration synchronization from the primary node ...
```

Workaround: Manually delete the **/root/temp** folder by using the `rm -rf /root/temp` command. This forces the **start.sh** script to rerun the initialization.

- If you have taken a backup of a Juniper Routing Director instance that includes the Active Assurance Victoria Metrics database (used for storing time-series data) and if you are restoring the instance, the GUI will fail to show the restored data. You might see a set of errors in the logs of the `metrics-service` in the `paa` Kubernetes namespace.

Workaround: Restart `paa-metrics` using the `kubectl rollout restart deployment paa-metrics -n paa` command

- In a multinode setup, taking a backup after a node has failed may cause the backup operation to fail.

Workaround: To perform a backup or restore operation, a fully operational setup with all services running is required. If a node is non-operational for any reason, we recommend that you resolve the issue before executing the backup or restore operation.

- When the cluster experiences a high load, some components, especially the Victoria Metrics Operator and ArangoDB Operator pods, may be restarted. This will not impact the cluster's functionality.

Workaround: None.

Resolved Issues

This section lists the issues resolved in Juniper Routing Director Release 2.8.0:

- When you download the Tunnels or Demands table using the **Download** option (**Planning > Networks > Offline Models > *Model-Name* > Open > Tunnels or Demands** tab), the CSV file incorrectly displays the Route Object value as *[object Object]*.

- Sometimes, continuous packet drops that happen under some specific trap codes may not be classified as a blackhole.

- If the node name contains uppercase letters, then the `request deployment backup schedule` command fails. This occurs because of the RFC 1123 host name validation rules.
- The signal functionality section (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Output Traffic**) does not display any alerts even when alerts exist.
- Management of Link Aggregation Groups (LAGs) from the Aggregated Ethernet Interfaces section of the topo resource instance is deprecated. Instead, LAGs are pre-created outside of the overlay service designs, which consumes only individual units on the Aggregated Ethernet (AE) interface. Additionally, upgrading to release 2.7.0 or later is not supported for instances using overlay-provisioned LAG interfaces.
- Alert notifications through webhook and e-mail do not contain certain information, such as which alert triggered the notification, the device ID, and so on. The email notification that is sent to the users does not include essential information, such as `message`, `hostname`, `deviceMAC`, `Key`, and so on.
- New network events (such as link status and LSP path changes) received through Web Socket do not support custom port forwarding. The Topology page (**Planning > Networks > Offline Models > *Model-Name* > Open**) does not automatically reflect network status changes.
- When you update the AS number on the Dynamic Topology tab of the Topology Settings Options page (**Observability > Topology > Settings icon**), the updated AS number is not reflected in the containerized routing protocol process (daemon) (cRPD).
- You cannot undelegate an SR LSP that is configured using NETCONF from Cisco devices.
- When you try to delete segment routing (SR) tunnels using the Routing Director GUI, the tunnels may not be deleted in the first attempt if there is a color conflict. This issue is seen in only SR tunnels that have a color conflict.
- During the installation of Test Agent on the JunOS EVO (such as ACX) devices, there can be few scenarios where the installation may get stuck in `CONFIGURING` state (due to slowness in router or some environment issue) for more than 10-15 minutes.
- Before you upgrade an EVPN service instance that was created in a release earlier to Release 2.5.0, ensure that you update the vpn-resources instance to Release 2.5.0.
- When a user with an administrator role clears the explicitly configured link speed on a physical interface within an EVPN LAG, the speed resets to 10 Mbps. Currently, only explicit speed values are supported. For help in removing the default value, contact Juniper Networks® Technical Assistance Center (JTAC).
- Although the *qinq* option is listed for the **Tag Type** field, we do not support *qinq* in VLAN-aware tagged EVPN services. We recommend that you do not select the *qinq* option as the L3VPN provisioning might fail with a validation error.

- The traffic statistics are not displayed on the Passive Assurance tab (**Orchestration > Instances > *Service-Instance-Name* > *Service-Instance-Name* Details**) for an L3VPN service with Integrated Routing and Bridging (IRB) interfaces.
- When you stitch a Layer 2 circuit with an L3VPN service, the Input Traffic Rate and Output Traffic Rate columns on the logical Interfaces accordion of the Passive Assurance tab (**Orchestration > Instances > *Service-Instance-Name* > *Service-Instance-Name* Details**) are blank.
- When there are two VLAN-aware tagged EVPN services, in placement section of an L3VPN service with Integrated Routing and Bridging (IRB) interfaces, both EVPN instances are not listed.
- When you change the IRB reference in the L3VPN configuration, the L3VPN service might display an error when you perform the placement operation.
- In a scaled setup, while provisioning VPN instances with a large Site Network Access (SNA) or access interfaces, you might notice VPN resource placement conflicts in the Arango database. Due to this conflict, you won't be able to onboard any device and provision a VPN instance.

  NOTE: For EVPN or L3VPN deployments, we recommend that you limit each VPN instance to a maximum of 20 SNAs or access interfaces.

- You might encounter a placement failure-related error when you try to upgrade the service instance for EVPN from Release 2.4.1 to Release 2.7.0.

  The placement fails because the speed field is marked as a key for *placement_interface* options in Release 2.7.0, while the options generated in the Release 2.4.0 service instance do not have the speed field.
- In a multihomed EVPN service, the GUI and the device configuration show different Ethernet Segment Identifier (ESI) values.
- After you upgrade Routing Director from Release 2.5.0 to 2.7.0, the provisioning of a service order with LAG interfaces configured fails. This is because provisioning of LAG Interfaces is deprecated starting from Release 2.7.0, and you need to configure LAG Interfaces through the network implementation plan.
- When you create an L3VPN service order with two OSPF sessions, and if one of those sessions fails, then the alerts are not displayed on the VPN State accordion of the L3VPN details for *L3VPN-Name* page (**Orchestration > Instances > Service Instances > *service-instance-name* hyperlink > Passive Assurance tab > L3VPN accordion > *L3VPN-Name***).
- The Input Traffic boundary is incorrectly displayed for Optical Rx power and Input Traffic KPIs. This issue occurs when you onboard a device and then stream data at a later point.
- The SPFRun anomaly might not get detected when both levels are enabled in IS-IS.

This issue does not apply to devices where only one level is enabled in IS-IS. A level can be disabled in IS-IS using the `set protocol isis level <level-number> disable` command.

- When onboarding and enabling BGP analytics for vMX routers, sometimes, the BMP configuration fails because Routing Director components cannot resolve the management IP address of the vMX routers. This issue is limited to vMX routers.