Release Notes

Published
2025-07-21

Juniper Routing Director Release 2.5.0

Software Highlights

Device Life-Cycle Management

- Device support:
  - **Basic support** (device management functions such as viewing inventory, system log and alarms, managing configuration using templates, and performing OS upgrades, orchestrating services using custom service definitions)—ACX710, ACX5048, ACX5096, EX2300, and SRX5400
  - **Full support** (device observability, routing observability, service orchestration, active assurance, network optimization, and viewing topology weather map)—ACX7020, MX104, MX2008, and MX2010,
- Configure aggregated Ethernet interface in a network implementation plan.

Observability

- Discover Shared Risk Link Groups (SRLGs) configured on the devices in your network.
- Track label-switched path (LSP) or SR policies events and changes made to LSP attributes.
- View performance graphs of custom KPIs.

Trust

- Monitor device integrity using the integrity accordion on the Trust dashboard.
- Run a compliance scan by configuring Trust settings.
- Update Trust-related definitions for SIRT, EOL, and PBNs in an air-gapped installation.

Service Orchestration

- Provision multihomed EVPN VPWS service with preconfigured LAG interfaces.
- Provision multihomed L2 circuit.
- Configure IRB interfaces for EVPN and L3VPN services.
- Configure Q-in-Q tunneling for EVPN and L3VPN services.

- Upload customized service designs.

Active Assurance

- Add and configure VLAN interfaces for Test Agent Appliances.
- Delete interfaces associated with Test Agent Appliances and Test Agent Applications.
- Configure IPv6 addresses in the management interfaces of Test Agent Appliances.

Network Optimization

- Set FAD constraints for tunnel profiles.
- Configure tunnel profiles using NETCONF.
- Implement tunnel diversity for enhanced network resilience.
- Configure tunnels using NETCONF.
- Set flexible algorithm definition (FAD) constraints for LSPs.
- Configure equal-cost multipath (ECMP) routing to distribute traffic and improve bandwidth utilization.
- Reroute LSPs based on link utilization threshold.

Planner

- Use the Planning menu on the Routing Director GUI to import your production network into Routing Director for offline modeling and simulation.

Administration

- Configure access control profiles to restrict access to Routing Director resources (devices, service instances, and tagging operations).
- Configure Routing Director to forward SNMP traps to an external system.
- Enforce security between the PCE (Path Computation Element) server and Path Computation Clients (PCC) by configuring a security mode.

Installation and Upgrade

- Deploy the Routing Director cluster on Ubuntu 22.04.05 kernel-based virtual machine (KVM) hypervisors.
- Schedule a periodic backup of Routing Director network configuration and telemetry information.

Table of Contents

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 7.

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 2.5.0.

- **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

This topic includes information on:

- Devices (Juniper Networks and third-party) supported by Routing Director
- Supported operating system (OS) versions
- 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 (device management functions, custom rules and custom service designs) • 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
MX Series	• MX104 (device management and observability functions only) • 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

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

Device Family	Device Series	Supported OS Version
EX Series	• EX3400 • EX2300 (Only device management functions, custom rules and custom service designs) • EX3400 • EX4300 [EX4300-32F (EMS only), EX4300-48mp] • 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
QFX Series	• QFX5110, QFX5120	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)	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 6 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 only device-controlled and delegated tunnels.

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

Vendor	Device Model	OS Version
Cisco Systems	Cisco 8202	IOS-XR 7.9.0
	Cisco NCS 57C3	IOS-XR 24.3.2
	Cisco NCS 5504	IOS-XR 7.3.2
	Cisco IOS XRv	IOS-XR 7.0.0
	Cisco ASR 9902	IOS-XR 7.0.0

gNMI Support

Table 3 on page 7 lists the gNMI support for Juniper Networks devices.

Table 3: 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 18.3R1 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

This section describes the features available in Juniper Routing Director Release 2.5.0.

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 2.5.0 extends device LCM to the following platforms and provides the following additional features:

- **Device Support**—You can onboard the devices listed in Table 4 on page 9 to Routing Director and manage them:

- **Table 4: Supported Devices and Functions**

Device	Supported Functions

ACX710, ACX5048, and ACX5096	• Basic device management functions:
	  • Viewing inventory, interface, system logs, and alarms
	  • Managing configurations using templates,
	  • Taking configuration backup
	  • Performing OS upgrades
	• Orchestrating services using custom rules and custom service definitions.

ACX7020 and ACX2200	• Device observability • Routing observability • Service orchestration, • Active assurance • Network optimization • Viewing topology weather map

EX2300	• Basic device management functions:
	  • Viewing inventory, interface, system logs, and alarms
	  • Managing configurations using templates,
	  • Taking configuration backup
	  • Performing OS upgrades
	• Orchestrating services using custom rules and custom service definitions.

Table 4: Supported Devices and Functions *(Continued)*

Device	Supported Functions
MX104, MX2008, MX2010, and MX2020	• MX104 (only basic device management functions). The following functions are supported for MX2008, MX2010, and MX2020 devices: • Device observability • Routing observability • Service orchestration, • Active assurance • Network optimization • Viewing topology weather map

SRX5400	• Basic device management functions:
	  • Viewing inventory, interface, system logs, and alarms
	  • Managing configurations using templates,
	  • Taking configuration backup
	  • Performing OS upgrades
	• Orchestrating services using custom rules and custom service definitions.

- **Configure aggregated Ethernet interface in a network implementation plan**—You can configure aggregated Ethernet interfaces in a network implementation plan so that:
  - LAGs are configured on a device during device onboarding.
  - The aggregated Ethernet interfaces are available as a resource for provisioning services soon after the device is onboarded.

  To use this feature, configure:

  - Maximum Transmission Unit (MTU) and tagging on the interfaces using port profiles in a network implementation plan.

- MTU, tagging, and LAG members in the network implementation plan.

[See Configure Port Profiles and Add a Network Implementation Plan].

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. Additionally, Routing Director provides a routing dashboard where you can actively monitor the overall routing health of your network in real time. The timely detection of anomalies enables you to take prompt action and minimize the impact of any issues.

Juniper Routing Director Release 2.5.0 provides the following additional observability features:

- **Discover SRLGs**—Use Routing Director to discover Shared Risk Link Groups (SRLGs) that you have configured on your network devices. SRLGs help you to identify the links that share the same risk factors and thereby support network redundancy and planning.

  You can view the list of SRLGs on the **SRLGs/Facilities** tab of the Topology page (**Observability > Network > Topology**).

  You need to establish a BGP Link State (BGP-LS) session between Routing Director and the device so that the device can advertise SRLG-related information configured in your network. Otherwise, you cannot view the list of SRLGs configured in your network on the SRLGs/Facilities tab.

  [See About the SRLGs Page.]

- **Track event history for LSPs**—Track changes made to label-switched path (LSP) attributes and events on the **Events** page of the **Tunnels** tab (**Observability > Network > Topology > Tunnels** tab). The graphical representation of bandwidth-related changes for an LSP on the Events page helps you to evaluate the changes in bandwidth over a period of time.

  In addition, you can use the **Show Path Changes** option to view highlights on the Topology map for any changes to the LSP. With this visibility, you can easily track and analyze routing changes.

  You can track event history only for PCC-initiated LSPs. By default, the retention policy for event history data is 3 days.

[See About the Tunnels Tab.]

- **View alerts and performance graphs of KPIs**—Routing Director generates alerts based on the KPI anomalies in your network. Use the KPI tab of the Monitor *Instance-Name* page (**Observability > Health > Custom KPI Collection > Rule Instantiation > *Instance-Name***) to view a graphical representation of the KPIs and alerts generated for all KPIs associated with a selected device and rule. You can also set triggers and threshold limits for an individual KPI when you define rule parameters.

  [See About the Custom KPI Collection Page.]

- **Predict events causing traffic loss using JRI**—Routing Director uses the Juniper Resiliency Interface (JRI) to collect routing, kernel, and forwarding exceptions that could potentially lead to traffic loss. These exceptions are displayed as Predictor Events in the Routing and MPLS accordion of the *Device-Name* page (Observability > Troubleshooting Devices > click *Device-Name*).

  Additionally, Routing Director offers a list of up to 10 events correlated to a predictor event and lists them in the decreasing order of their relevance to the predictor event.

  [See Predictor Events.]

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 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.

Juniper Routing Director Release 2.5.0 provides the following additional trust and compliance features:

- **Monitor device integrity**—You can monitor the device integrity using the integrity accordion on the Trust tab **(Observability > Health > Health Dashboard > Trust)**. On this accordion, you can view:

  - Percentage of healthy devices and the total number of unhealthy devices based on integrity of devices.
  - A graphical representation of average health of all devices.
  - KPIs such as Hardware EOL and Software EOL that affect the overall network health.

  You can use this information to perform corrective actions. You can also click **View Details** to view detailed information about device health and KPIs that affect the integrity of devices.

[See About the Integrity Health Page.]

- **Run a compliance scan by configuring trust settings**—Run a compliance scan by configuring trust settings on the Trust tab (**Observability > Health > Health Dashboard > Trust**) of the Health Dashboard. Select a trust plan, benchmark document, profile, and tailoring document from the Trust Settings page to obtain a desired level of compliance. The results of the scan are displayed on the Compliance accordion.

  [See About the Trust tab.]

- **Update trust-related definitions for SIRT, EOL, and PBNs in an air-gapped installation**—Update all trust-related KPIs when Routing Director does not have access to the Internet, using the `request paragon trust data update hostname` command.

  View the updated KPIs on the Trust tab of the Health Dashboard (**Observability > Health > Health Dashboard > Trust**). You can use this information to detect anomalies and perform corrective actions,.

  [See Update Trust KPIs in an Air-Gapped Installation.]

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 2.5.0 provides the following additional service orchestration features:

- **Configure an IRB interface for EVPN and L3VPN services**—You can configure an integrated routing and bridging (IRB) interface to simultaneously support Layer 2 (L2) bridging and Layer 3 (L3) routing on the same interface for EVPN and L3VPN services. To configure an IRB interface, you must:

  - For EVPN services—Enable **Layer3 Gateway** and select the placement options for the IRB interface on the Add Connection page (**Orchestration > Instances > + > E-LAN EVPN CSM > Customer Site Settings > Site Network Access > + > Add Connection**).

    Configuring an IRB interface for EVPN services is qualified only for the following scenarios:

- EVPN with Ethernet interface port mode.
- EVPN with Ethernet interface VLAN mode.

- For L3VPN services—Select IRB as the connection type and configure the virtual gateway IP address and MAC address for the IRB interface on the Add Connection page (**Orchestration > Instances > + > L3VPN > Customer Site Settings > Site Network Access > + > Add Connection**). You can also select the placement options for the IRB interface.

  Configuring an IRB interface for L3VPN services is qualified only for the following scenarios:

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

  You can also configure an IRB interface when you add node details to the topology resource pool and devices in the network implementation plan.

  [See Add EVPN or EVPN-VPWS Service Site Details, Add L3VPN Service Site Details, Add a Network Implementation Plan, and Configure Resource Pools for Resource Instances.]

- **Provision a multihomed EVPN-VPWS service with preconfigured LAG interfaces**—You can configure multihoming on your Ethernet VPN–virtual private wireless service (EVPN-VPWS) service. To configure multihoming you must have pre-provisioned the link aggregation group (LAG) aggregated Ethernet interfaces on the device. You can provision LAG interfaces by logging in to the device and configuring the interfaces manually or through the Routing Director GUI by using network implementation plans. After provisioning the LAG interface on the device, add the interface under the Access Interfaces section (**Modify *Resource-Instance* > Pop/Site > + > Node > + > Access Interfaces > + > Access**). Then set the **Lag** toggle button to **True** when configuring a topology resource instance for provisioning an EVPN-VPWS service.

  [See Create a Topology Resource Pool and Add a Network Implementation Plan.]

- **Provision a multihomed L2 circuit**—Routing Director supports configuring multihomed Layer 2 (L2) circuits in the following two modes:

  - Backup mode—In the default backup mode, a single active pseudowire exists between the access node and the remote PE node. The backup pseudowire is not pre-signaled. If the active pseudowire fails, the backup pseudowire becomes active after a connection is established.
  - Hot-standby mode—In this mode, the backup pseudowire is pre-signaled. If the active pseudowire fails, the backup pseudowire becomes active immediately. As the backup pseudowire is pre-signaled, the switchover time to the backup pseudowire is reduced, minimizing traffic disruption.

To configure multihoming, you must add three devices to the L2 circuit.

[See Add L2 Circuit VPN Nodes.]

- **Configure Q-in-Q tunneling for EVPN and L3VPN services**—You can assign a Q-in-Q tag, in addition to a Dot1q tag, to an Ethernet interface for E-LAN EVPN CSM and Layer 3 (L3VPN) services. Assign both an inner customer VLAN ID and an outer service VLAN ID to the same interface when you use a Q-in-Q tag.

  You can assign Q-in-Q as the tag type for interfaces when you add node details to the topology resource pool and devices in the network implementation plan.

  [See Add EVPN or EVPN-VPWS Service Site Details, Add L3VPN Service Site Details, Configure Resource Pools for Resource Instances, and Add a Network Implementation Plan.]

- **Assign tags to service instances and network implementation plans**—You can assign tags to logically group and categorize service instances and network implementation plans.

  To assign tags to a service instance or network implementation plan, click **More > Manage Tags** on the Service Instances and Network Implementation Plan pages respectively.

Network Optimization

The network optimization use case in Routing Director 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) through an intent-based approach.

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, each path in a tunnel must be configured and provisioned 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 2.5.0 provides the following network optimization features:

- **Implement tunnel diversity for enhanced network resilience**—You can pair label-switched paths (LSPs) so that they do not overlap at any point between their ingress (starting point) and egress (ending

point). Tunnel diversity is crucial for network resilience, as the feature ensures that if one path fails, the other remains operational.

We support three types of tunnel diversity:

- Link diversity—Ensures LSPs do not share common links. If one LSP uses a specific link, the other LSP in the pair uses a different link, thus avoiding shared points of failure.
- Shared Risk Link Group (SRLG) diversity—Ensures LSPs avoid links that share the same risk factors.
- Site diversity—Ensures LSPs avoid the same nodes and therefore the same sites. This tunnel diversity type includes both link and SRLG diversity features.

Use the Add Diverse Tunnels page (**Observability > Network > Topology > Tunnels tab > Provisioning > Diverse Tunnels**) to implement tunnel diversity in your network.

[See Add Diverse Tunnels.]

- **Support for ECMP routing**—Routing Director supports equal-cost multipath (ECMP) routing to distribute traffic and improve bandwidth utilization on multiple links to the same destination. In the PathFinder Settings section of the Organization Settings page (**Settings Menu > System Settings > Organization Settings**), select one of the following placement methods for the **ECMP Placement Method** field:
  - Random—Randomly selects one of the ECMP paths
  - Least filled—Selects the path with the maximum available bandwidth
  - Most filled—Selects the path with the minimum available bandwidth

  Routing Director evaluates link metrics and applies the specified placement method if multiple ECMP paths are available.

  [See Reroute LSPs.]

- **Reroute LSPs based on link utilization threshold**—Routing Director automatically reroutes a label-switched path (LSP) when the utilization of a link exceeds a specified threshold, ensuring optimal network performance and preventing congestion.

  For automatic rerouting of LSPs, you must configure the following parameters in the PathFinder Settings section of the Organization Settings page (**Settings Menu > System Settings > Organization Settings**):

  - Link Utilization Threshold—Specify a threshold value for link utilization (in percentage) so that Routing Director can reroute LSPs when the traffic on a link exceeds this value.
  - Threshold Rerouting Interval—Specify the minimum interval (in minutes) after which Routing Director reacts to any threshold violations.

[See Reroute LSPs.]

- **Configure secondary or standby LSPs**—Add a secondary or standby label-switched path (LSP) to a primary LSP so that an alternate route is chosen when the primary route fails.

  The tunnel ID, source node, destination node, and IP address of a secondary or standby LSP are identical to that of the primary LSP. However, secondary and standby LSPs have the following differences:

  - A secondary LSP is not signaled until the primary LSP fails.
  - A standby LSP is signaled regardless of the primary LSP status.

  To add a secondary or standby LSP, select a tunnel on the Tunnels tab of the Topology page (**Observability > Network > Topology**) and click **Provisioning > Secondary/Standby Tunnel**.

  [See About the Tunnels Tab.]

- **Add or remove LSP delegation**—Use the Configure LSP Delegation page (**Observability > Network > Topology > Tunnel tab > Delegation > Configure Delegation**) to delegate the management of Path Computation Client-controlled LSPs (PCC-controlled LSPs) to the Path Computation Element (PCE). After delegation, these LSPs are managed by the PCE and are classified as PCE-delegated LSPs. You can also return the control of these LSPs to the PCC by removing the delegation, and these LSPs are then reclassified as PCC-controlled LSPs.

  [See Add and Remove LSP Delegation.]

- **Configure tunnels and tunnel profiles using NETCONF**—Use NETCONF in addition to Path Computation Element Protocol (PCEP) to configure tunnels. Routing Director provides the NETCONF option in the Provisioning Method field when you configure a tunnel (**Observability > Network > Topology > Tunnels tab > Provisioning drop-down list > Tunnel**).

  In addition, NETCONF is supported when you create tunnel profiles using the intent-based network optimization method. Routing Director provides the NETCONF option in the Provisioning Method field when you configure tunnel profiles (**Observability > Network > Topology > Tunnels tab > Provisioning drop-down list > Tunnel**).

  When NETCONF is the provisioning method, the tunnel is statically provisioned and the associated configuration statements is part of the router configuration file. This tunnel is considered as a device-controlled tunnel.

  [See Add a Tunnel and Add a Tunnel Profile.]

- **Set FAD constraints for LSPs and Tunnel Profiles**—You can apply Flexible Algorithm Definition (FAD) constraints to your LSP so that LSPs comply with the defined FAD constraints, maintaining optimal network paths. If a compliant path cannot be found, the LSPs remain unprovisioned.

On the Add Tunnel page (**Observability > Network > Topology > Provisioning > Tunnels tab > +**), you can select the Flexible Algorithm (flex algo) IDs.

In addition, you can apply FAD constraints to tunnel profiles so that LSPs that are created using the intent-based network optimization method comply with the defined FAD constraints. On the Add a Tunnel Profile page (**Network Optimization > Network > Network Optimization Intent > Tunnel Profiles tab > +**), you can select the Flexible Algorithm (flex algo) ID to associate the constraint to the LSP.

You can assign flex algo IDs only for segment routing (SR) and Segment Routing for IPv6 (SRv6) LSPs.

[See Add a Tunnel and Add a Tunnel Profile.]

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, new capacity or services, and so on.

Planner generates a topology view of a network, enabling the addition, removal, and reconfiguration of 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 2.5.0 provides the following planner feature:

- **Analyze and simulate offline networks**—Use the Planning menu on the Routing Director GUI to import either from Routing Director or using a JSON document. You can then perform an exhaustive failure analysis of all single link, node or SRLG failures, and generate reports to understand the critical failure events that could adversely impact network traffic.

  [See Network Planner Overview and Network Planner Workflow.]

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 2.5.0 provides the following additional Active Assurance features:

- **Add VLAN interfaces on Test Agent Appliances**— Add and configure VLAN interfaces from the **Interfaces** tab of the ***Test-Agent-Name*** (**Inventory > Active Assurance > Test Agents**) page. You can configure VLAN-specific parameters such as the parent interface, VLAN ID, namespace, isolation parameter, and so on. You can add interfaces only on Test Agent Appliances. In addition, you can also set the interface as the management interface or Network Time Protocol (NTP) source for the Test Agent Appliance.

  [See About the Test Agent-Name Page.]

- **Delete interfaces on Test Agents**—You can delete interfaces associated with Test Agent Appliances and Test Agent Applications from the **Interfaces** tab of the ***Test-Agent-Name***page (**Inventory > Active Assurance > Test Agents**).

  You can mark non-orphaned interfaces associated with Test Agents for deletion by using the **Delete** option and then commit the changes, and you can directly delete orphaned interfaces on Test Agents by using the **Delete** option. You can also unmark an interface for deletion by using the **Undelete** option before committing the configuration. This action restores the interface.

  [See About the Test Agent-Name Page.]

- **IPv6 support for management interfaces**—In addition to IPv4 addresses, you can configure IPv6 addresses in the management interface of a Test Agent Appliance. You can assign various types of IPv6 addresses, such as static addresses and addresses assigned using DHCP or SLAAC. You can also enable IPv6 support for the NTP client. With this support, a Test Agent Appliance can synchronize time by using an IPv6-based NTP client.

  [See Install Test Agent Appliance.]

Administration

Routing Director Release 2.5.0 provides the following administration features to manage users, sites, and organizations:

- **Enable PCEP security**—Enforce security between the PCE (Path Computation Element) server and Path Computation Clients (PCC) by configuring a security mode to use certificates. Routing Director supports the following PCEP security modes:

  - `auto-detect`—Accepts connection regardless of whether PCEP security is enabled.

- `strict-disable`—Accepts only non-secured PCEP connections. This is the default mode.
- `strict-enable`—Enforces security on all PCEP sessions. If a device does not have security enabled, its connection is rejected.

To ensure that communication between the PCE server and its clients is encrypted and secure, you must set the security mode to `strict-enable` or `auto-detect` using Paragon Shell.

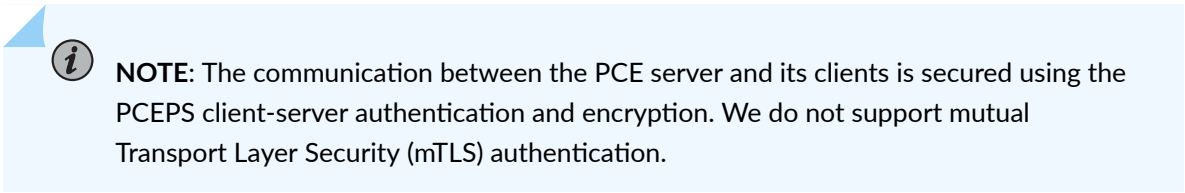

NOTE: The communication between the PCE server and its clients is secured using the PCEPS client-server authentication and encryption. We do not support mutual Transport Layer Security (mTLS) authentication.

Routing Director supports both custom certificates and system-generated certificates. If you are using a system-generated certificate, the server certificate (tls.crt) and private key (tls.key) are automatically generated; otherwise, you must upload a custom certificate.

[See Enable PCEP Security.]

- **Use access control profiles to restrict access to Routing Director resources**—With access control profiles, superusers can enforce conditions for granting Routing Director resource access to network administrators. Access controls are implemented by using tags.

  The following operations are access controlled:

  - Device operations: Device update, reset, reboot, deletion, and so on.
  - Service instance operations: Service instance update, force sync, placement update, deprovisioning, and so on.
  - Tagging operations: Tag assignment, update, and deletion.

  A network administrator cannot access a resource if:

  - A resource is not assigned a tag.
  - The user is not assigned an access control profile that include access to that resource.

  [See Access Control Profile Overview.]

- **Configure Routing Director to forward SNMP traps to an external system**—Routing Director can receive SNMP traps (SNMPV2c and SNMPV3) from Juniper devices and forward them to an external system for fault management.

  To use this feature, you must configure SNMP and SNMP forwarding endpoints in Organization Settings (**System Settings > Settings Menu) under SNMP Configuration and SNMP Forwarding Endpoints respectively.**.

[See Configure SNMP.]

Juniper Routing Director Installation

Juniper Routing Director Release 2.5.0 provides the following installation-related features:

- **Support for Ubuntu 22.04.05 KVM hypervisors**—You can deploy the Routing Director cluster on Ubuntu 22.04.05 kernel-based virtual machine (KVM) hypervisors, in addition to Red Hat Enterprise Linux (RHEL) 8.10 KVM hypervisors.

  [See Create the Node VMs.]

- **Schedule periodic backup**—Schedule a periodic backup of your Routing Director network configuration and telemetry information. You can save the backed-up information on the cluster node or upload the information to a remote storage location. You can also view and delete scheduled backups.

  [See Schedule Backups.]

Beta Features

Juniper Routing Director Release 2.5.0 provides Beta support for the following features:

- **Routing Director Chatbot**— Use Routing Director chatbot (LLM Connector) to facilitate the use of natural language to query network status and obtain troubleshooting information, without using CLI commands.

  LLM Connector can help you:

  - Retrieve device information.
  - Execute Junos OS operational commands.
  - Save data (configuration and logs) to a file.
  - Retrieve a list of all VPNs in your network and their details, metrics, and health information.
  - Fetch information about customers and service instances associated with customers.
  - Get insights based on the telemetry collected from the device.

- Plot various metrics and state graphs. For example, LLM Connector can plot data related to CPU usage, device temperature, memory usage, packet loss, BGP session status, and so on.
- Fetch insights based on the telemetry data derived from the metrics collected by the Test Agents.

To use the LLM Connector tool, you must set up a large language model (LLM). The recommended LLM model for LLM Connector is GPT-4 and GPT-4o.

[See LLM Connector Overview.]

- **Monitor device health and temperature using AI-ML**—Routing Director uses artificial intelligence-machine learning (AI-ML) to monitor key performance indicators (KPIs) for device health to detect anomalies. Routing Director monitors the KPIs for the following components:
  - Fans [Revolutions per minute (RPM)]
  - Line cards (CPU utilization, memory utilization, and temperature)
  - Routing Engine (CPU utilization, memory utilization, and temperature)
  - Interfaces (optical temperature, optical Tx power, optical Rx power, input traffic, and output traffic)

  Routing Director also performs root cause analysis (RCA) of device temperature anomalies.

  [See Automatically Monitor Device Health and Detect Anomalies.]

- **Detect blackholes using AI-ML**—Routing Director uses artificial intelligence–machine learning (AI-ML) to detect blackhole (packet drops) on a device and generates a Blackhole Detected alert to notify about the event. On the Traffic Loss page (Observability > Troubleshoot Devices > click *Device-Name* > Routing and MPLS accordion > Traffic Loss), you can view:
  - Graphs—Depicts input packet rate, output packet rate, and packet drop.
  - Alerts—Provides details of alerts and traffic flows affected by packet drops.

  [See Detect Blackholes.]

- **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.]

- **Support for Test Agent Appliances**—Use Routing Director to register and run Test Agent Appliances in your network. A Test Agent Appliance is a full-fledged Test Agent with a built-in operating system, which provides you full control over network configuration and supports advanced functionalities. The Test Agent Appliance is based on the Debian Linux operating system. It is delivered as:

  - Dedicated Test Agent—Download the Test Agent Appliance software image and install the software image on custom x86 hardware.

  - Test Agentvirtual network function—Upload a Test Agent software image to a virtualization platform and run the platform as a virtual machine (VM) on a hypervisor.

  After you install a Test Agent Appliance on a platform, Routing Director discovers the Test Agent Appliance. You can view the discovered Test Agent Appliance on the Test Agents **(Inventory > Active Assurance > Test Agents)** page.

  [See About the Test Agents Page.]

- **Install Test Agent Appliance on custom x86 hardware**—You can install a Test Agent Appliance on x86 hardware. For the installation, create a bootable USB flash drive with the Test Agent Appliance software image and use the USB flash drive to boot the hardware. Register the installed Test Agent Appliance with Routing Director to start collecting measurements in your network.

  As an alternative to this deployment method, you can use virtual machine-based deployment to install a Test Agent Appliance. Routing Director supports installation of Test Agent Appliance on Proxmox.

  [See Install Test Agent Appliance.]

- **Configure RPM Cisco TWAMP Reflector using GUI**—You can configure an RPM Cisco TWAMP Reflector measurement testing from the Measurement Designer page (**Observability > Active Assurance > Measurement Designer > Create a Test/Monitor**).

  The RPM Cisco TWAMP Reflector task measures network performance by running Two-Way Active measurement protocol (TWAMP) tests on Cisco devices. This test sets up a reflector for TWAMP senders to evaluate metrics like round-trip delay, jitter, and packet loss between TWAMP clients and Cisco devices (reflectors).

  [See Tests and Monitors Overview.]

Known Issues

IN THIS SECTION

This section lists the known issues in Juniper Routing Director.

Device Life-Cycle Management

- In a scaled environment, the network implementation plan provisioning might sporadically fail if multiple provisioning jobs are executed consecutively at a high rate. You might see the following error in the workflow GUI:

  ```
  Provisioning is failing at upload-network-resources due to the error "numbered is in use and cannot be deleted; app.ErrorCode:3"
  ```

  Workaround: Retry publishing the failed network implementation plan, and space out the provisioning jobs to avoid this issue.

- In clusters with inter-node delays in the tens of milliseconds, the workflow execution time increases. As a result, the time required to onboard devices and instantiate services also increases.

  Workaround: None.

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

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

- If you upgrade a network implementation plan to the latest service design version and publish it without making changes, the service design is upgraded again instead of publishing the network implementation plan.

  Workaround: Instead of directly publishing a network implementation plan after service design upgrade, ensure that you edit the network implementation plan before publishing.

- After you upgrade a service design, you must manually upgrade the service design that is used in the existing network implementation plans.

  To upgrade the infrastructure service design used by an implementation plan:

  1. On the Network Implementation Plan page (**Inventory > Device Onboarding > Network Implementation Plan**), select the plan and click **More > Upgrade Service Design**.

     The Upgrade Service Design page appears.

  2. On the Upgrade Service Design dialog box, click the drop-down list to select the service design version to which you want to upgrade the network implementation plan, and click Yes.

  A message indicating that the service design upgrade is in process is displayed. After a successful upgrade, the Service Design column displays the new version of the infrastructure service design used by the network implementation plan.

- If a network implementation plan includes multiple devices, the device onboarding might not start automatically for some devices. However, you might notice that the device status is displayed as Connected on the Inventory page (**Inventory > Devices > Network Inventory**).

  Workaround: You must manually trigger the onboarding process for such devices. To manually trigger the onboarding process:

  1. On the Network Implementation Plan page (**Inventory > Device Onboarding**), select the network implementation plan and click **More > Trigger Onboarding**. The Trigger Onboarding page appears.

  2. Click the **Show All devices** option and then select the device for which you want to restart onboarding.

  3. Click **OK**. A message indicating that onboarding has started appears.

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

- Upgrading from release 2.4.0 to release 2.5.0 fails for all components in routing observability.

  Workaround: Delete all databases related to routing observability and reinstall routing observability components afresh.

- Upgrading from release 2.4.0 to release 2.5.0 for fails for routing observability BGP components due to differences in implementations meant to improve ingest performance.

  Workaround: Delete all BGP related databases and install release 2.5.0 BGP components.

- Under scale ingest, the service bgp-measurements-compute may restart continuously after crashing. This affects the calculation of aggregates such as the total routes graph in **Observability > Routing > Routing Explorer > Routing Status**. Other related dashboards and tables are unaffected.

  Workaround: None.

- After you upgrade to Release 2.5.0, you might notice that the following pages might take more than 10 minutes to display device-related information:

  1. Inventory page (**Observability > Troubleshoot devices >** ***Device-Name***)
  2. Devices tab on the Topology page **(Observability > Topology**)

  Workaround: After you upgrade, restart services using the `kubectl rollout restart -n streams deployment/papi-mon-v2` command.

- The Alert icon and the alert message are out of sync during the first event change on the Custom KPI Collection page (**Observability > Health > Custom KPI Collection**).

  Workaround: Wait for the next data refresh cycle, which automatically corrects the alert icon display issue.

- On the LLM Connector (Routing Director chatbot), you might not be able to continue the conversation if the session expires or if you open a new conversation

  Workaround: Execute the following commands on the Linux root shell of a primary node:

```
# kubectl get svc -n common | grep opensearch-cluster-master

# curl -X POST :9200/ask-paragon-chat-sessions/_rollover/
```

- 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.

- During heavy ingest scenarios such as onboarding of routers for the first time or router maintenance windows, it takes sometime for the total number of routes to be reflected on the Routing Status graph (**Observability > Routing > Routing Explorer Routing Status** tab).

  If there are any events in the network, the Routing Status graph or the Routing Updates table (**Observability > Routing > Route Explorer > Routing Updates**) might display the data with substantial latency. We expect that the latency is reasonable during steady state operation of the network.

  Also, the statistics in the Device tab (**Observability > Routing > Route Explorer > Routing Status**) or in the Adjacencies tab (**Observability > Routing > Route Explorer**) are updated with low latency (1 to 5 minutes).

  Workaround: None.

- 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.

- 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 5 on page 28.

Table 5: 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 5: 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 31 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 31 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 6 on page 31 for more information.

  Workaround: None.

  Table 6: 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 6: 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

- In a scaled environment when you provision multiple VPN instances parallelly, some of the instances may fail with the following error in the workflow:

  ```
  "Failed to lock service instance for uploading: Service order is locked. Please try again
  later."
  ```

  Workaround: Retry provisioning for the failed VPN instances.

- The following error message is displayed, although the service design upgrade is successful.

  Service Design upgrade failed for 1 instance(s)

  The incorrect error message is displayed only when you try to upgrade the service design in a VPWS instance that was originally provisioned in Release 2.4.1 and is now upgraded to Release 2.5.0.

  Workaround: None.

- On a scaled setup, the Resource Instance page takes a longer time to load due to slower responses of instances and the */service-designs* REST API calls.

  Workaround: None.

- Before you upgrade an L2VPN 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.

Workaround: None.

- In a scale set up, the View Network Resources page (**Orchestration > Service > Resource Instances > More**) might be unresponsive.

  Workaround: Do one of the following:

  - Select the network resources and click the edit option to view and download the details in the JSON format.

    or

  - Use the REST API, `/service-orchestration/api/v1/orgs/{org-id}/placement/network-elements`, to view the details.

- When creating l2-addr resource instance, adding only the LACP Admin Key without System ID results in a failure.

  Workaround: When you create an l2-addr resource instance with LACP Admin Key, also create an entry in the System ID table to avoid failure.

- If more than 2600 L3VPN services are provisioned, the placement service intermittently suffers a break in service which may cause service placement to fail. The placement service will automatically recover from this failure.

  Workaround: Re-provision any services that failed placement due to this issue.

- The Service Designs page (**Orchestration > Service Catalog**) might list certain internal service files. You can ignore these files.

  Workaround: None.

- The Add Access Parameter page displays all options for the **Tag Type** field (Untagged, dot1q, or qinq), irrespective of whether you have chosen tagged or untagged interfaces for a customer edge (CE) device.

  Workaround: None.

- When you modify the tag type from **dot1q** to **qinq** for the existing CE devices (that is, the **ce** or **ces** option of the **CE Spec** field) on the Access page (**Orchestration > Resource Instances > Modify *Resource-Instance-Name* > *Resource-Instance-Name* > +** icon above the Access Interfaces table), the VLAN information displays both tag types resources (numbered and dual numbered).

  Workaround: Create a topology resource based on the tag type and do not modify the topology resources based on the tag type for the existing CE devices.

- If different L3VPN services are running on the same IFD using different MTU values, then service provisioning fails.

Workaround: Ensure that the MTU values are the same for L3VPN services that share the same IFD.

- 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.

- The device name is not displayed when you hover over the **View Details** hyperlink in the Relevant Events section of the L3VPN accordion (**Orchestration > Instances > Service Instances > *Service-Instance-Name* hyperlink > *Service-Instance-Name* Details > Passive Assurance tab**).

  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 designs in bulk.

  Workaround: We recommend that you upgrade only one service design at a time.

Active Assurance

- After you perform the **Undelete** operation, the commit operation for the Test Agent fails. This issue occurs regardless of the interface involved.

  Workaround: Delete the orphan VLAN interfaces. After deletion, the commit functionality is restored for the affected Test Agent.

- 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.

- When you click the Distribution tab on the *Application-Name* page (**Observability > Health > Health Dashboard > Active Assurance (Tab) > Applications (Accordion) > View Details**), the page hangs and you might not be able to see metrics and site-related data for a Measurement.

  Workaround: None.

- 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

- You cannot configure **OSPF TE Metric A** and **OSPF TE Metric Z** parameters for a link that has OSPF enabled on it.

  Workaround: None.

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

  Workaround: None.

- The link utilization reroute threshold feature works only if the same threshold values are specified on the nodes of a link.

  Ensure that you specify the same values for **Util Reroute Threshold AZ** and **Util Reroute Threshold ZA** fields on the Edit Link page (**Observability > Network > Topology > Link tab > *Link-Name* > Edit**).

- If the configuration database is locked exclusively by a root terminal session, then tunnel provisioning fails and the status is displayed as *Unknown*.

  Workaround: Use the Edit LSP option in the GUI and re-provision the tunnel.

- 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.

- 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).

  Workaround: In addition to updating the AS number using the Routing Director GUI, you must log in to the cRPD CLI and update the AS number.

- In rare cases, the topology-related information might be lost or incomplete. This issue occurs due to the inaccessibility of databases.

  Workaround: Restart the toposerver to restore topology information.

  To restart the toposerver connect to the primary node of the Routing Director cluster and run the `kubectl -n $(kubectl get namespaces -o jsonpath='{.items}' | jq -r '.[] | select(.metadata.name| startswith("pf-"))|.metadata.name') rollout restart deployment toposerver` command.

- 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

- Planner reports are deleted based on the retention policy settings. However, reports older than the retention period are deleted only when the scripts are run. The cleanup scripts are scheduled to run at midnight everyday.

  Workaround: None.

Trust

There are no known issues in this release.

Administration

There are no known issues in this release.

Installation and Upgrade

- Sometimes when taking a backup, the OpenSearch backup status is Not Available but others are successful. For example:

```
Backup (config data) status for backup ID 20250709-100554:
    Paragon configuration : SUCCESS
Backup (telemetry data) status for backup ID 20250709-100554:
    Timescale job status : SUCCESS
    Opensearch job status : NOT AVAILABLE
    Insights Victoriametrics job status : SUCCESS
    Pathfinder Victoriametrics job status : SUCCESS
```

  This is due to the status flag not being updated.

  Workaround: Validate the OpenSearch backup status using the following commands:

1. Determine the backup-server service IP address from the Linux root shell.

```
# kubectl -n common get svc backup-server -o jsonpath='{.spec.clusterIP}'
```

2. Retrieve the status of OpenSearch backup.

```
# curl -s http://backup-server service IP:9000/snapshot/telemetry/opensearch/common/backup ID | jq '.message'
```

The status should be displayed as `SUCCESS`.

- After restoring from a backup, sometimes Airflow installation fails when you try to reinstall all application services using the `request paragon service start` command. It is possible that the Airflow pods are still shutting down.

  Workaround: Perform the following steps:

  1. Stop all running services.

```
> request paragon service destroy
```

  2. Run the following command from the Linux root shell to remove all Airflow db-pooler pods.

```
# kubectl -n airflow scale deploy --replicas=0 atom-db-pooler
```

  3. Wait till all the pods terminate. Use the following command to ensure that all Airflow db-pooler pods are removed and no pods are in `Running` status. The command output must be empty.

```
# kubectl -n airflow get pods -l application=db-connection-pooler
```

  4. Perform the restore operation as usual.

  5. After restore is complete, if the db-pooler pods don't come up automatically. Start the db-pooler pods manually.

```
# kubectl -n airflow scale deploy --replicas=2 atom-db-pooler
```

- You might encounter the following error while upgrading or redeploying Routing Director:

```
This host is not master node 1, where this Paragon cluster was initially installed from!
```

  This issue occurs because **/root/epic/host.ip** is empty.

  Workaround: Manually re-populate **/root/epic/host.ip** with the local IP address before upgrading or redeploying Routing Director. You can also optionally make the **/root/epic/host.ip** file immutable to prevent overwrites.

Resolved Issues

This section lists the issues resolved in Juniper Routing Director Release 2.5.0:

- If you have upgraded the topology resource from Release 2.2.0 or Release 2.3.0 to Release 2.4.0 and if you later edit and provision a service instance (L3VPN or EVPN) that was created in an older release (Release 2.3.0 or Release 2.2.0), then the provisioning of the service instance fails.
- While creating or modifying an EVPN service order, you cannot configure multiple VLAN IDs on the Aggregated Ethernet (AE) interface. The EVPN considers the AE port as a single resource and therefore an AE interface cannot be reused across service instances even when the VLAN IDs on the AE IFL differ.
- When the worker node is down, there might be issues if you create an organization or onboard a device.
- When you run the `request paragon deploy cluster` or `request paragon service start` commands, sometimes deployment may fail because the **config.yml** is empty. In such cases, the log file might display an error similar to this:

```
usage: ansible-playbook [-h] [--version] [-v] [--private-key PRIVATE_KEY_FILE]
                        [-u REMOTE_USER] [-c CONNECTION] [-T TIMEOUT]
                        [--ssh-common-args SSH_COMMON_ARGS]
                        [--sftp-extra-args SFTP_EXTRA_ARGS]
                        [--scp-extra-args SCP_EXTRA_ARGS]
                        [--ssh-extra-args SSH_EXTRA_ARGS]
                        [-k | --connection-password-file CONNECTION_PASSWORD_FILE]
                        [--force-handlers] [--flush-cache] [-b]
                        [--become-method BECOME_METHOD]
                        [--become-user BECOME_USER]
```

```
                        [-K | --become-password-file BECOME_PASSWORD_FILE]
                        [-t TAGS] [--skip-tags SKIP_TAGS] [-C]
                        [--syntax-check] [-D] [-i INVENTORY] [--list-hosts]
                        [-l SUBSET] [-e EXTRA_VARS] [--vault-id VAULT_IDS]
                        [--ask-vault-password | --vault-password-file VAULT_PASSWORD_FILES]
                        [-f FORKS] [-M MODULE_PATH] [--list-tasks]
                        [--list-tags] [--step] [--start-at-task START_AT_TASK]
                        playbook [playbook ...]

Runs Ansible playbooks, executing the defined tasks on the targeted hosts.


< output snipped >


  --become-method BECOME_METHOD
                        privilege escalation method to use (default=sudo), use
                        `ansible-doc -t become -l` to list valid choices.
  --become-user BECOME_USER
                        run operations as this user (default=root)
  -b, --become          run operations with become (does not imply password
                        prompting)
```

- Due to the changes in telemetry paths, you cannot view IS-IS data for ACX7020 devices on the Routing and MPLS accordion (**Observability > Health > Troubleshoot > Devices > *Device-Name***).
- If a device is discovered through a BGP-LS peering session even before you onboard the device, then duplicate LSPs are created when a PCEP session is established with the device. In rare cases, the duplicate LSPs that are created will continue to remain.
- The "vpn_svc_type" service type is displayed as "pbb-evpn" instead of "evpn-mpls" on the Routing Director GUI and through the REST API.
- You cannot run multiple versions of plug-ins on a Test Agent.
- If no valid interface option is available for a CE and PE device combination, then the Interface drop-down will be empty.
- The Route Explorer page (**Observability > Routing**) displays data only if you have installed Junos OS or Junos OS Evolved Release 23.2 or earlier.
- In a rare case of running multi-level ISIS protocols on a link, the topology map might not be updated or might not reflect the latest live operation status.

- After you update a Monitor or a Test Template that is created by another user, the **Updated By** column on the Monitors (**Observability > Active Assurance**) and Test Template (**Inventory > Active Assurance**) pages do not reflect the name of the user that modified the Monitor or the Test Template.
- When you onboard devices in batches, due to Kubernetes' horizontal pod autoscaling of airflow-worker pods, the onboarding might fail for the devices that are in the middle of the onboarding process.
- The Output Traffic rate column on the Logical Interface accordion (**Orchestration > Instances > Service Instances page > *service-instance-name* hyperlink > *Service-Instance-Name* Details**) displays some data even when there is no traffic through the devices.
- While modifying an existing L3VPN service instance, if you try to remove a device that is already a part of a network implementation plan then the modify workflow fails.
- After you upgrade Paragon Automation from Release 2.2.0 to Release 2.4.0, ensure that you upgrade the L3VPN service instance before you upgrade the topology resource instance; otherwise, you might encounter issues.
- If you upgrade Paragon Automation from Release 2.3.0 to 2.4.0, then you might not be able to modify VLANs for Site Network Accesses on the existing L3VPN service instances.
- The maximum size of a configuration template supported is 1 MB and not 10 MB as indicated in the error message on the GUI.
- The devices table on the Devices tab (**Observability > Health > Health Dashboard > Active Assurance (Tab) > Click any accordion > View Details > Affected Items** tab) does not list devices that have unhealthy measurements.
- When you click a Monitor on the Monitors page (**Observability > Active Assurance**), the *Monitor-Name* page takes approximately a minute to load the data. This issue occurs only when there are more number of events in the system.
- If you had installed Test Agent on a router while using Juniper Paragon Automation Release 2.3.0 or earlier releases, and later if you upgrade to Paragon Automation Release 2.4.0, then Test Agent that was installed earlier will not be upgraded to the latest version of Test Agent. You cannot run Tests or Monitors on this router.
- Due to changes in XML Path Language (XPath), some custom rules cannot collect KPI information from the device.
- The View Network Resources page (**Inventory > Device Onboarding > Network Implementation Plan > More**) doesn't display AE interfaces-related details.
- Onboarding a QFX device to Routing Director fails if **Trust** is enabled in a device profile applied to the QFX device.

- You might not be able to view the Tests page (**Observability > Active Assurance**) if your role type is Observer.
- The *vmrestore* tool restores data into vmstorage pods. While performing restore, the tool creates a lock file that prevents any other application from accessing data during the restore phase. However, sometimes the vmrestore tool fails to clear the lock file and the vmstorage pods cannot access data.
- If you have onboarded a Cisco device, but later changed the TLS settings on the device (either turn it on or off), the status of the device will show as Disconnected on the Inventory page.