

Release Notes

Published
2025-11-02

Juniper Paragon Automation Release 2.3.0

Software Highlights

- **Device Life-Cycle Management**—Select the IP version (IPv4 or IPv6) to be used in the outbound SSH command generated for device onboarding, and import brownfield devices to a network implementation plan.

Onboard and manage Juniper Networks® EX4300 Series switches.

- **Observability**—Provision a Labeled Switched Path (LSP), view historical data for links and tunnels, and customize the topology map to view real-time information on delays and bandwidth utilization for links or tunnels.
- **Trust and Compliance**—Use the Trust dashboard to view the overall trust score, compliance status of devices, and monitor vulnerabilities that affect targets in the network in real-time.
- **Service Orchestration**—Manage allocation of interface placement options for a customer to provision an L2 circuit service.
- **Active Assurance**—Use the Active Assurance tab on the Health Dashboard page to monitor the health of various measurements.

Evaluate the QoS in your network using Y.1731 Delay Measurement (DM) and Y.1731 Synthetic Loss Measurement (SLM) native plug-ins.

- **Administration**—Bulk upgrade device images.
- **Installation and Upgrade**—Gracefully shut down or reboot a node or cluster using Paragon Shell CLI.

Table of Contents

Introduction | 1

Licensing | 2

Supported Junos OS Releases, Devices, and Browsers | 3

Installation and Upgrade | 5

New Features | 6

Deprecated Features | 13

Known Issues | 13

Resolved Issues | 30

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® Paragon Automation 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.

Paragon Automation is based on a modern microservices architecture with open APIs. Paragon Automation is designed with an easy to use UI that provides a superior operational and user experience. For example, Paragon Automation 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.

Paragon Automation takes a use case-based approach to network operations. When you execute a use case, Paragon Automation 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.

Paragon Automation supports the following use cases:

- **Device life-cycle management (LCM)**—Allows you to onboard, provision, and then manage a device. Paragon Automation 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, Paragon Automation notifies you about network issues using alerts, alarms and events, which you can use to troubleshoot issues affecting your network.
- **Trust and compliance**—Automatically checks whether the device complies with the rules defined in the Center for Internet Security (CIS) benchmarks document. In addition, Paragon Automation 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.

For details about these use cases and other features of Paragon Automation, see ["New Features" on page 6](#).

In summary, Paragon Automation 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 Paragon Automation.

Licensing

To use Paragon Automation and its features, you need:

- **Product Entitlement**—To use Paragon Automation and its use cases.



NOTE: Product entitlements are honor-based and not enforced for Paragon Automation Release 2.3.0.

- **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 Paragon Automation GUI. For more information, see [Manage Device Licenses](#).

Supported Junos OS Releases, Devices, and Browsers

[Table 1 on page 3](#) lists the supported Junos OS release, devices, and browsers in Juniper Paragon Automation.

Table 1: Supported Junos OS release, devices, and browsers

Supported Junos OS	• Junos OS Evolved releases 23.2R2, 22.4R2, 22.2R3, 23.4R2, and 24.2R1. • Junos OS releases 23.2R2, 22.4R2, 22.2R3, 23.4R2, and 24.2R1.
--------------------	---

Supported Juniper Devices

- ACX7024
- ACX7024-X
- ACX7100-32C
- ACX7100-48L
- ACX7348
- ACX7332
- ACX7509
- PTX10001-36MR
- PTX10002-36QDD
- PTX10004
- PTX10008
- PTX10016
- MX204
- MX240
- MX304
- MX480
- MX960
- MX10003
- MX10004
- MX10008
- vMX
- EX4300

Supported third-party devices	• Cisco Network Convergence System 57C3 (Cisco NCS57C3) • Cisco Network Convergence System 5504 (Cisco NCS5504) • Cisco 8202 Router • Cisco IOS XRv Router • Cisco Aggregation Services Routers 9902 (Cisco ASR9902) NOTE: Only EMS and service provisioning-related features are supported for these third-party devices.
Supported Browsers	The latest version of Google Chrome, Mozilla Firefox, and Safari.

Installation and Upgrade

Release 2.3.0 is no longer available for download from the [software download](#) site. We recommend that you install release 2.4.1 or upgrade to release 2.4.1.

To install Juniper Paragon Automation Release 2.4.1 afresh, download the **paragon-2.4.1.9371.ova** file from the Juniper Paragon Automation [software download](#) site. Perform the steps described in the *Installation and Upgrade Guide* to install release 2.4.1 and log in to the Web GUI. See [Install Paragon Automation](#) for information.

If you have already installed Juniper Paragon Automation Release 2.3.0, upgrade to release 2.4.1 by downloading the **upgrade_paragon-release-2.4.1.9371.g4407b53b58.tgz** file available on the [software download](#) site. See [Upgrade Paragon Automation](#) for information.

New Features

IN THIS SECTION

- [Device Life-Cycle Management | 6](#)
- [Observability | 7](#)
- [Trust and Compliance | 8](#)
- [Service Orchestration | 9](#)
- [Active Assurance | 10](#)
- [Administration | 11](#)
- [Juniper Paragon Automation Installation | 11](#)
- [Beta Features | 12](#)



NOTE: Juniper Paragon Automation Release 2.3.0 is an early adopter release. In case you encounter any issues, you can contact Juniper Networks Technical Assistance Center (JTAC).

This section describes the features available in Juniper Paragon Automation Release 2.3.0.

Device Life-Cycle Management

Device life-cycle management (LCM) encompasses the entire lifecycle of the device, from installing the device onsite, bringing the device under management, monitoring the device when it is in production, and finally decommissioning the device.

Juniper Paragon Automation Release 2.3.0 extends device life-cycle management to the following platform and provides the following additional features:

- **Support for EX Series Switches**—In addition to routers, you can onboard Juniper Networks® EX4300 Series Switches to Paragon Automation.

In this release of Paragon Automation, you can view the details of the supported EX Series Switches on the Inventory page (**Inventory > Devices > Network Inventory**) and manage them from the Troubleshoot Devices page (**Observability > Health > Troubleshoot Devices**).

Among the EX4300 device models, only the EX4300-48MP model supports telemetry. For the complete list of devices supported in this release, see "[Supported Junos OS Releases, Devices, and Browsers](#)" on page 3.

- **Import brownfield devices to a network implementation plan**—Paragon Automation enables you to import brownfield devices to a network implementation plan. Paragon Automation populates the device information in the plan. You can use advanced monitoring and configuration features of Paragon Automation such as automatic software upgrade to manage the brownfield devices.

[See [Add a Network Implementation Plan](#).]

- **Select IPv4 or IPv6 addresses for device onboarding**— Paragon Automation provides the **Select IP Version** drop-down list (**Inventory > Network Inventory > Add Devices**) from which you can select the IP version (IPv4 or IPv6) to be used in the outbound SSH command generated for device onboarding.

IPv4 is the default IP version that is used in the outbound SSH command.

[See [Add a Juniper Device](#).]

Observability

Paragon Automation enables you to view your entire network topology in real time, monitor network health, be notified of any anomalies in the network, and also get guidance on the remediation of these anomalies. With observability, Paragon Automation monitors and analyzes the network and its components by using key performance indicators (KPIs), device logs, and metrics, and notifies you about network issues through alerts and alarms. Additionally, Paragon Automation runs connectivity tests using synthetic traffic to identify connection issues between devices in your network. The timely detection of anomalies enables you to take prompt action and minimize the impact of any issues that occur.

Juniper Paragon Automation Release 2.3.0 provides the following additional observability features:

- **Provision a labeled -switched path (LSP)**—Paragon Automation enables you to provision LSPs (tunnels) by using Path Computation Element Protocol (PCEP).

You can create a tunnel with provisioning type as PCEP on the Tunnels tab of the Topology (**Observability > Topology**) page. The Path Computation Element (PCE) creates and provisions a PCE-initiated tunnel.



NOTE: The PCE-triggered secondary path is not supported for Cisco devices.

You can create RSVP, SRV6, and segment-routed LSPs, and view the LSPs on the topology map.

[See [Add a Tunnel](#).]

- **View historical data for links and tunnels**—Paragon Automation enables you to view historical data for links and tunnels. The historical data provides you valuable insights on past trends and patterns.

On the Topology page (**Observability > Topology**), you can view historical data for:

- Interface traffic for a link
- Interface delay for a link
- Tunnel traffic

[See [About the Link Tab](#) and [About the Tunnels Tab](#).]

- **Enhanced topology visualization**—On the Topology page (**Observability > Topology**), you can customize the topology map to view real-time information about delays and bandwidth utilization for links or tunnels. Labels on the map indicate the delays and the bandwidth utilization data.

In addition, you can click the legend on the bottom-right corner of the topology map to select the color to be displayed for a certain percentage of link utilization. Link utilization is the last measured interface traffic pushed between devices.

[See [View Network Topology Details](#).]

Trust and Compliance

Paragon Automation 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 devices and device components. Paragon Automation distills the outcomes of these checks into a single trust score that you can use to determine how trustworthy a device is.

Juniper Paragon Automation Release 2.3.0 provides the following additional trust and compliance features:

- **Monitor network trust**— The Trust tab on the Health Dashboard (**Observability > Health > Health Dashboard > Trust**) displays the overall trust score, device compliance, and vulnerabilities that affect targets in the network. You can use this information to perform corrective actions.

The Trust tab consists of:

- **Trust pane**—View the trust score of the targets, trust plan applied to the network, overall trust trend, and a graph of the average trust score for the past 30 days.
- **Vulnerabilities accordion**—View the percentage of healthy devices and the total number of unhealthy devices based on device vulnerabilities, and a graph that displays the average health of all devices. You can also view KPIs such as proactive bug notifications (PBNs) and advisories that affect overall network health.
- **Compliance accordion**—View the percentage of healthy devices and the total number of unhealthy devices based on compliance of devices. The compliance of devices is measured against the KPIs provided in the benchmark document. You can also view a graph that displays the average health of all devices.

Click **View Details** on any accordion to view the related health pages where you can view the percentage of healthy devices and a list of KPIs. You can also view the total number of unhealthy devices that are grouped by OS version and device model.

[See [About the Trust Tab.](#)]

Service Orchestration

Service orchestration is the process of designing, configuring, validating, deploying, and monitoring a network service. Paragon Automation automates the entire life cycle of a network service by providing workflows that execute the tasks to be completed to deliver a service. You can provision various network services by using predefined service designs written in YANG. 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, Paragon Automation activates the automated workflow to provision the service in the network. After provisioning, Paragon Automation monitors the service by automatically setting up Juniper® Paragon Insights and Juniper® Paragon Active Assurance instances to monitor network health and measure service quality.

Juniper Paragon Automation Release 2.3.0 provides the following additional service orchestration feature:

- **Automatically list interfaces placement options for an L2 circuit service instance**—Use the **Update Placements** button (on the **Orchestration > Instances > Add/Modify E-Line L2Circuit NSM > VPN Nodes** page) to automatically list all tagged and untagged interfaces placement options available for a customer to provision a Layer 2 (L2) circuit service. Select and allocate the provider edge (PE) device

interfaces from the listed options on the Configure Network Access page (**Orchestration > Instances > Modify/Add E-Line L2Circuit NSM > VPN Nodes > VPN Node > Network Accesses > Configure Network Access**) when you create or modify a service instance.

[See [Add L2 Circuit VPN Nodes](#).]

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 the Test Agents at strategic locations in your network and install them on routers running Junos OS Evolved, x86 hardware, or on virtual machines (VMs). Paragon Automation uses RPM to collect metric data for Juniper Networks® MX Series Universal Routers and Juniper Networks® PTX Series Routers.

Juniper Paragon Automation Release 2.3.0 provides the following additional Active Assurance features:

- **Support for additional native plug-ins**—Paragon Automation enables you to evaluate the QoS in your network using the following plug-ins:
 - **Y.1731 Delay Measurement (DM)**—Measures network performance on the parameter of latency during a transmission. Y.1731 DM is important for services that need low latency (such as video conferencing or VoIP).
 - **Y.1731 Synthetic Loss Measurement (SLM)**—Measures network performance on the parameter of loss of frames during a transmission. This parameter is important for maintaining QoS and SLA guarantees.

[See [Tests and Monitors Overview](#).]

- **Monitor the health of various measurements**—The Active Assurance tab on the Health Dashboard page (**Observability > Health > Health Dashboard > Active Assurance**) enables you to monitor the health of various measurements by using the following accordions:
 - **Applications**—View the overall health of measurements related to HTTP, DNS, and RPM HTTP plug-ins in your network.
 - **L3 Network**—View the overall health of measurements related to TCP, UDP, TWAMP, TWAMP reflector, ping, RPM UDP, RPM TCP, and RPM TWAMP plug-ins in your network.

- **L2 Network**—View the overall health of measurements related to Y.1731 SLM and Y.1731 DM plug-ins in your network.
- **TV & Multicast**—View the overall health of measurements related to IPTV, OTT, and Netflix speedtest plug-ins in your network.

[See [About the Active Assurance Tab.](#)]

Administration

Paragon Automation Release 2.3.0 provides the following administration features to manage users, sites, and organizations:

- **Bulk upgrade device images**—You can now upgrade the image installed on multiple devices of a particular model at the same time.
- **Tag key naming updates**—You can start the name of a tag key by using an uppercase or lowercase letter. Earlier, you could start the name of a tag key by using only a lowercase letter. The tag key can have periods (.) and hyphens (-) in addition to letters, underscores (_), and numbers. The maximum number of characters allowed is increased to 200.

[See [Upgrade Images on Devices.](#)]

[See [Manage Tags.](#)]

Juniper Paragon Automation Installation

Juniper Paragon Automation Release 2.3.0 provides the following installation-related features:

- **Gracefully shut down or reboot a node or cluster**—Paragon Shell enables you to gracefully shut down a node virtual machine (VM) or the whole Paragon Automation cluster using the request `paragon shutdown type (node | cluster)` command. You can also reboot a node VM or the whole cluster using the request `paragon reboot type (node | cluster)` command.

[See [Shut Down and Reboot Nodes.](#)]

Beta Features

Juniper Paragon Automation Release 2.3.0 provides Beta support for the following features:

- **Manage allocation of network resources for service provisioning**—Placement is the process of allocating network resources for provisioning services. Use the **Update Placements** button to automatically assign all the placement options available for a customer to provision L3VPN and EVPN services. After placement options are assigned, you can select and allocate network resources such as provider edge (PE) devices and interfaces, VLANs, and so on from the available options to provision services for the customer.

Use the **Reset Placements** button on the Summary page of the service creation wizard to delete all allocated placement options for all sites in a service instance. Use the **Reset Placements** button on the Edit Connection page to delete all allocated placement options for an individual site.

[See [Manage Placement Configurations for Service Instances.](#)]

- **Configure IPv6 IP connection and access diversity parameters**—You can configure IPv6 addresses, in addition to IPv4 addresses, when you configure IP connection parameters and static route, OSPF, and BGP routing protocols under Layer 3 VPN (L3VPN) Site Network Access.

[See [Add L3VPN Service Site Details.](#)]

- **Paragon Automation Chatbot**—Paragon Automation provides the LLM Connector tool to facilitate the use of natural language to query network status and obtain troubleshooting information, without the need for 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.

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

[See [LLM Connector Overview.](#)]

Deprecated Features

The following features are deprecated in Juniper Paragon Automation Release 2.3.0.

- The Network Score page (**Trust > General > Network Score**) that provides a dashboard of real-time information about the trustworthiness of your network is not available in this release. Alternatively, you can use the Trust tab (**Observability > Health > Health Dashboard > Trust Tab**) of the Health Dashboard to view this information.
- The Snapshots page (**Trust > General > Network Score > Snapshots**) that enables you to view a snapshot of a target's performance over a period of time is not available in this release. Alternatively, you can use the Trust tab (**Observability > Health > Health Dashboard > Trust Tab**) of the Health Dashboard to view snapshots.
- You cannot access the Snapshots page by clicking **Trust > General > Network Score > Snapshots** in this release. However, you can click the **Score** link in the Identity & Location accordion (**Observability > Health > Troubleshoot Devices > Device > Device-Name > Identity & Location**) to view snapshots.

Known Issues

IN THIS SECTION

- [Device Life-Cycle Management | 14](#)
- [Observability | 16](#)
- [Service Orchestration | 22](#)
- [Active Assurance | 26](#)
- [Trust | 27](#)
- [Administration | 27](#)
- [Installation and Upgrade | 28](#)

This section lists the known issues in Juniper Paragon Automation.

Device Life-Cycle Management

- If the device onboarding fails, the device onboarding status is displayed as *Status is not available* in the Devices section of the Network Implementation Plan page (**Inventory > Device Onboarding > Network Implementation Plan**).

Workaround: For such devices, initiate the outbound SSH connection on the router to restart the onboarding workflow.

- In this release, single node failure and restart might cause inconsistencies in device inventory.

Workaround: None.

- Paragon Automation 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.

- If you have enabled the Trust option in the device profile, the onboarding workflow may fail with the following error:

Onboarding workflow failed reason: task_failure, Trust score computation failed.

Workaround: Retry the onboarding process after a few minutes.

- Onboarding an ACX7024 device fails with the following error:

task_failure, Failed to launch Software update. Timed out or waiting for the launch message. Try again later.

Workaround: Restart the onboarding process using the service orchestration cMGD CLI. Perform the following steps:

1. Log into the primary node using SSH.
2. Exit out of the default paragon CLI to the Linux root shell.
3. Log in to the service orchestration CMGD CLI.

```
kubectl exec -it -n foghorn deployment/cmgd -- bash  
cli
```

4. Get the organization UUID from the GUI. To find your organization's UUID, go to **System Settings > Organization Settings** on the Paragon Automation UI banner and then copy the organization's UUID.

5. Configure the organization UUID in the cMGD CLI:

```
configure
set foghorn:core org-id org-uuid
commit and-quit
```

6. Retrieve the service order.

```
show service order status customer-id_instance-id
request service project add /data-models/projects/network-resource.tgz
request service project add /data-models/projects/onboard.tgz
request service order sync customer-id_instance-id
```

7. Reset the node state.

```
configure
delete foghorn:customers customer customer-id services instance-id infrastructure
infrastructure-ntw network-nodes network-node <node-name> onboard
set foghorn:customers customer customer-id services instance-id infrastructure
infrastructure-ntw network-nodes network-node <node-name> node-type paper-node
commit and-quit
request service order load merge customer-id_instance-id
```

8. Log in to the router.
9. Restart the onboarding by restarting the connection.

```
configure
deactivate system services outbound-ssh
commit
activate system services outbound-ssh
commit and-quit
```

- The onboarding of a device fails if you configure only IPv6 addresses at the time of installation and enable **Trust** in the device profile. This issue does not occur if you have configured IPv4 addresses.

Workaround: To successfully onboard the device, do one of the following:

- Disable **Trust** in the device profile and retry onboarding. In this case, the Trust compliance score will not be calculated.
- At the time of installation, include IPv4 addresses in addition to IPv6 addresses.
- The View Network Resources page (**Inventory > Device Onboarding > Network Implementation Plan > More**) doesn't display AE interfaces-related details.

Workaround: You can view the AE interfaces-related details in the **View active config** link of the Configuration accordion (**Observability > Troubleshoot Devices > Device-Name**).

Observability

- The Hardware accordion does not display the following information for the listed devices:
 - Chassis temperature (chassis-temperature) for MX204, MX240, MX304, MX10004, MX10008, and MX10016 devices.
 - Charts related to the fan speed (rpm-percent) for MX480, MX960, MX10004, MX10008, and MX10016 devices.
 - Power supply module temperature (psm-temperature) for MX204, MX480, and MX960 devices.
 - CPU utilization (cpu-utilization) information is not available for PTX10002 device.
 - Line card charts are not available for some ACX Series, MX Series, and PTX Series devices as the flexible PIC concentrator (FPC) fields are not supported on these devices. See [Table 2 on page 16](#) for more information.

Table 2: Line Card Charts Support

Device Family	Device Series	FPC Fields Not Supported
ACX Series	ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7509, ACX7348	fpc-temperature, fpc-cpu-utilization, fpc-buffer-memory-utilization

Table 2: Line Card Charts Support (*Continued*)

Device Family	Device Series	FPC Fields Not Supported
MX Series	MX204, MX240, MX304, MX480, MX960, MX10004, MX10008, MX10016	fpc-temperature, fpc-cpu-utilization
PTX Series	PTX10002	fpc-temperature, fpc-cpu-utilization, fpc-buffer-memory-utilization

- On the Interfaces accordion, forward error correction (FEC) corrected errors and 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 graphs related to cyclic redundancy check (CRC) errors display *No Results Found* as the CRC errors-related data is not streamed from the devices for the management ports.

Workaround: None.

- 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 some rare cases, the duplicate LSPs may remain.

Workaround: If you see duplicate LSPs, restart the EdgeAdapter pod.

- After the primary node is switched off, the OC-term and GNMI state on the Remote Management accordion (**Observability > Troubleshoot Devices > Device-Name**) are displayed as disconnected.

Workaround: You can do one of the following:

- Offboard and onboard the devices, or
- Restart the OC-term pod.
- Sometimes, the graph on the Output Traffic Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > Device-Name > Overview > Interfaces (accordion) > Output Traffic data**

link) displays a sudden spike in data. This issue occurs because some devices might erroneously send data more than once with the same time stamp.

Workaround: None.

- 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 Paragon Automation 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 19](#) 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 19](#) 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 Paragon Automation 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.

- The values for Input Utilization and Output Utilization are represented in terms of percentage of interface utilization instead of Mbps.

This issue is seen in Physical Interfaces Details for *Service-Instance-Name* tab (**Orchestration > Instances > Service Instances > *Service-Instance-Name* hyperlink > *Service-Instance-Name* Details > Passive Assurance**) and Input Traffic Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > *Device-Name* > Overview > Interfaces (accordion) > Input Traffic data-link**).

Workaround: To get the actual traffic rate value, you should multiply the values displayed by interface speed.

- Not all optics modules support all the optics-related KPIs. See [Table 3 on page 20](#) for more information.

Workaround: None.

Table 3: 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
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.
 - On the Interfaces Details for *Device-Name* Page (**Interfaces accordion > Interfaces data-link**), the link flaps count erroneously drops to zero and reverts to the same count as before.

Workaround: None.

- For EX9204 devices, the Optical Tx power and Optical Rx Power graphs on the Pluggables Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > Device-Name > Overview > Interfaces accordion > Pluggables data-link**) do not show any warnings for critical thresholds.

Workaround: None.

- For EX4300-48MP devices, the Signal Functionality graph on the Input Traffic Details for *Device-Name* page (**Observability > Health > Troubleshoot Devices > Device-Name > Overview > Interfaces accordion > Pluggables**), does not display any data.

Workaround: For Pluggables (Optical Interfaces), we recommend that you install Junos OS Release 24.4R2 on your device.

- While creating or editing an LSP, the **Routing Path** field (**Add or Edit Tunnel > Path** tab) for a node may not list IPv6 addresses. However, you should be able to select any Link or Node to create the path.

Workaround: None.

- You won't be able to provision an LSP using the Paragon Automation GUI if the provisioning type is SRv6.

Workaround: We recommend that you provision SRv6 LSPs using the REST APIs.

- The changes are not reflected when you update the following fields on the Constraints tab of the Edit LSP page:
 - Admin Group Include All
 - Admin Group Include Any
 - Admin Group Exclude

This issue occurs only when you try to edit an LSP that doesn't have at least one of these attributes set at the time of LSP creation.

Workaround: None.

- If you have installed a Junos OS or Junos OS Evolved release other than Release 22.1R1.10, you may not be able to view the historical data for link delays. For any other versions of Junos OS or Junos EVO release, you may see discrepancies in the graph as delay statistics are sent only when there are changes.

Workaround: None.

- Your hard disk may be full because of the LSP event-related data retention.

Workaround: If you have enabled the PCEP protocol, we recommend that you monitor the disk usage and clean up the unused LSP events-related data in the Postgres database. To do the cleanup:

1. Check the disk usage utilization of the Postgres database.
 - a. SSH to one of the primary nodes.
 - b. Type `exit` to exit from Paragon Shell.
 - c. Run the `paragon-local-volume-check` command.
 - d. Look for the size of `pgdata-atom-db*` and see if it has grown over time.

2. Clean up the unused LSP event-related data, in case of highly-utilized disk due to Postgres usage.

- a. SSH to one of the primary nodes
- b. Type `exit` to exit from Paragon Shell.
- c. Run the `paragon-db` command.
- d. Run the `\c ns_pcs;` command.
- e. Run the `truncate table pcs_lsp_event;` command.

If a read-only error occurs, you can choose a different primary node, depending on the role you have for connecting to Postgres,

- The historical data for tunnel traffic (**Observability > Topology > Tunnels > View > Tunnel Traffic**) is not displayed for ACX Series devices.

Workaround: None.

Service Orchestration

- The "vpn_svc_type" service type is displayed as "pbb-evpn" instead of "evpn-mpls" on the Paragon Automation GUI and through the REST API.

Workaround: None.

- Sometimes, the apply insights configuration (**appy_insights-config**) fails if you try to provision a service without properly deleting a previously provisioned service or a device.

For example, if you release the router without off-boarding or deleting a service, then the apply insights configuration fails when the same service or device is used in another organization.

Workaround:

- If there are stale services and devices, run the following REST APIs from the cMGD container of the **foghorn** namespace to delete stale services and devices, and rerun the workflow:
 - `curl --request DELETE <http://config-server.healthbot:9000/api/v2/config/services/device-group/<device-group> name>/`
 - `curl --request DELETE <http://config-server.healthbot:9000/api/v2/config/device-group/<device-group> name>/`
 - `curl --request POST <http://config-server.healthbot:9000/api/v2/config/configuration>/`

- If there are stale network-groups, run the following REST APIs from the cMGD container of the **foghorn** namespace to delete the stale network-groups, and rerun the workflow:
 - `curl --request DELETE <http://config-server.healthbot:9000/api/v2/config/services/device-group/<network-group> name>/`
 - `curl --request DELETE <http://config-server.healthbot:9000/api/v2/config/device-group/<network-group> name>/`
 - `curl --request POST <http://config-server.healthbot:9000/api/v2/config/configuration>/`
- The EVPN service order creation fails if you try to create an EVPN service order by importing an existing JavaScript Object Notation (JSON) file.

Workaround: If you are using a JSON file, ensure that you clear the placement section before you publish the service order.

- For some devices such as ACX7204, if you configure VLANs on unused ports, the following error occurs:

VLAN must be specified on tagged interfaces.

Workaround: This issue is caused by the default factory configuration on the port. Delete the default factory configuration on the ports that you plan to use.

- 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.

- Although multiple VLAN IDs are available in the topology resources, the Placement section of the EVPN service order lists only one VLAN ID in the drop down.

Workaround: To fix this issue:

1. Edit the EVPN service order to add new VLAN IDs. You can add the VLAN IDs under the Tagged Interface section.
 2. Clear the Placement section by deselecting the device name.
 3. Save and publish the service order.
- While modifying a service order, you cannot clear the existing placements.

Workaround: If publishing the service order fails due to existing placements, you can export the failed service order to JSON format and then create a new service order or modify an existing service order by importing this JSON file. During the importing process, delete the placements and then publish the service order.

- 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.

Workaround: None.

- VLAN drop down under Placement section doesn't display all the available VLANs as per the topology service order and it shows only the selected VLAN.

Workaround: None.

- During device onboarding, pings from the device to Microsoft Azure and Google Cloud Platform endpoints fail.

Workaround: None.

- While creating an EVPN service order, if the MAC Address Limit that you have specified is out of the defined range then the service order fails.

Workaround: Specify a value that is within the defined range and then republish the service order.

- While creating or modifying an EVPN service order, the MAC Address Limit configuration is ignored if you specify Drop as the action to be taken when the upper limit for customer MAC addresses exceed.

Workaround: None.

- Since you can manage resources from within network implementation plans, a blank topology instance (*topo*) is auto created on the Resources Instances page (**Orchestration > Service > Resource Instances**). This topology instance is a read-only instance and is owned by the network operator.

Workaround: None.

- If you have referenced a device in a service order and if you try to release the router from the organization, then you cannot delete the service.

Workaround: We recommend that you deprovision the service and offboard the device before you release the router from the Network Inventory page.

- 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.

- After you upgrade Paragon Automation from Release 2.1.0 to Release 2.2.0, L2VPN and L3VPN accordions (**Orchestration > Instances > Service Instances > *Service-Instance-Name* hyperlink > *Service-Instance-Name Details* > Passive Assurance tab**) show no data for service instances that were provisioned in Release 2.1.0. This issue does not occur if you install Paragon Automation afresh.

Workaround: You need to recreate the VPN service after you upgrade to Juniper Paragon Automation Release 2.2.0.

- 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.

Workaround: On the Monitors page, stop all the monitors associated with the device that has to be deleted in the service instance. After the relevant monitors are stopped, you can proceed with modifying the L3VPN service instance.

- Scheduling provisioning of service orders does not work if you upgrade Paragon Automation to Release 2.2.0.

Workaround: None.

- The Order History tab on the *L3VPN-Name* Details page (**Orchestration > Instances > *Service-Instance-Name* hyperlink**) lists all the order history even if the service instance is deprovisioned and provisioned again.

Workaround: None.

- When you upgrade an L2 circuit service design, you are prompted to update the L2 address even when there is no dependency on L2 addresses.

Workaround: Before you upgrade an L2 circuit service design that is associated with a service instance, you must upgrade the L2 address, VPN resources, and topology services using the Paragon Automation GUI.

- While creating topo resources, the node name is not populated. This issue occurs only after you upgrade to Release 2.3.0.

Workaround: Perform the following steps to ensure that the node name is populated:

1. Navigate to the Troubleshoot Devices page (**Observability > Health**) or Network Inventory (**Inventory > Devices**) page.
2. Select one or more devices for which the node name is not populated and click **More > Assign to Site**.
3. On the Assign Devices to Sites page, select a site other than the site the device was originally assigned to.

4. After you assign a new site, repeat Step 1 through Step 3 and assign the device to the original site.

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

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

- 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.

Workaround: None.

- After you upgrade Paragon Automation to Release 2.3.0 and upgrade the L3VPN service instance to the latest service design, you may not see the physical interface of one of the devices.

Workaround: Restart the AED pods by running the following command:

```
kubect1 delete pod -n healthbot analytical-engine-pod-names
```

Ensure that the pod names are separated by a space.

To get the list of analytical-engine pods names, run the following command:

```
kubect1 get pods -n healthbot | grep analytical-engine
```

- You cannot provision an L3VPN service on an ACX7024 device with Junos OS Release 23.2R2 installed.

Workaround: None.

Active Assurance

- 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.

- You cannot run multiple versions of plug-ins on a Test Agent.

Workaround: When you upgrade Paragon Automation, restart all measurements before creating any new 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.

Workaround: None.

- The streams are not generated when you create a Test with a DNS plug-in and the following event is raised:

Could not get nameserver from resolv.conf

This issue occurs when the Test is associated with a Test Agent that runs on a Juniper Networks router with Junos OS EVO installed, and you don't specify the Name Server field while configuring a Test.

Workaround: Ensure that you specify a value for the Name Server field while configuring a Test.

- 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.

Workaround: None.

- 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.

- 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.

Workaround: None.

Trust

There are no known issues in this release.

Administration

There are no known issues in this release.

Installation and Upgrade



WARNING: The backup and restore functionality available in Paragon Shell is inconsistent and might not work as intended. We recommend that you use the snapshot feature available on VMware clients to back up and restore your cluster. For more information, see [Back Up and Restore Paragon Automation](#).

- The backup and restore functionality has the following caveats:
 - You cannot restore data backed up from a setup with a different release of Paragon Automation to a setup with the current release.
 - Paragon Automation backs up only application configurations such as devices, sites, service orders, and so on. Since a backup does not store the certificates and infrastructure services configurations, that information must be kept unchanged during restoration.
 - Resources allocated to the network won't be preserved after a restore and you must ensure that you release the allocated resources during the window between taking a backup and performing a restore.

Workaround: None.

- If a node in the cluster is not operational, the status of the vector pod from the node that is not operational is displayed as *Running*, even though the node status is reported as *Not Ready*. This is due to an existing Kubernetes issue. See <https://github.com/kubernetes/kubernetes/issues/117769>.

Workaround: You can do one of the following:

- Monitor the metric, *kube_daemonset_status_number_ready*. When the value for this metric drops to 3, you can manually check from which vector the data is missing.
- Set a query and an alert for the *kube_daemonset_status_number_ready* metric in Grafana.
- You might encounter RKE2-related issues if you change the hostname after you set up a cluster.

We recommend that you do not change the hostname after a cluster is set up.

Workaround: None.

- When the worker node is down, there might be issues if you create an organization or onboard a device.

Workaround: Do not create an organization or onboard a device when a worker node is down. You must wait until the cluster recovers and then create an organization or onboard a device. Recovered

state is when all the pods are either in *Running* or *Pending* states and are not in any intermediate states like *Terminating*, *CrashloopbackOff*, and so on.

- If you have powered off one of the primary nodes, you may not be able to log in to the Juniper Paragon Automation GUI.

Workaround: Restart papi-ws using the following Paragon Shell CLI command:

```
request paragon cluster pods reset service papi-ws namespace papi operation restart
```

- Upgrading from Release 2.1.0 to Release 2.2.0 on VMs configured with minimum required specifications might fail due to a failed airflow-worker instance. Verify the reason for failure.

```
root@pa1:~# kubectl get pods -A | grep -v Running
NAMESPACE
NAME                                READY   STATUS
RESTARTS      AGE
airflow
worker-6cc5484f95-78cfv             0/1     Pending   0          30m
healthbot
cgvzz                                0/2     Completed 0          2d11h
kube-system
zs8pg                                0/1     Completed 0          2d11h
kube-system
qzbnw                                0/1     Completed 1          2d11h

<output snipped>

rook-ceph
cds6v                                0/1     Completed 0          6h38m
root@pa1:~# kubectl describe pods airflow-worker-6cc5484f95-78cfv -n airflow
Name:          airflow-worker-6cc5484f95-78cfv
Namespace:     airflow
Priority:       0

<output snipped>

Tolerations:      node.kubernetes.io/not-ready:NoExecute op=Exists for 30s
                  node.kubernetes.io/unreachable:NoExecute op=Exists for 30s

Events:
  Type    Reason             Age          From          Message
  ----    -
Warning  FailedScheduling  21m         default-scheduler  0/4 nodes are available:
4 Insufficient cpu. preemption: 0/4 nodes are available: 4 No preemption victims found for
```

```
incoming pod.
Warning FailedScheduling 5m10s (x4 over 20m) default-scheduler 0/4 nodes are available:
4 Insufficient cpu. preemption: 0/4 nodes are available: 4 No preemption victims found for
incoming pod.
root@pa1:~#
```

Workaround: Perform the following steps:

1. Lower the following CPU resource requests from the Linux root shell.

```
# kubectl patch deployment -n mems mems --type "json" -p '[{"op":"add","path":"/spec/template/spec/containers/0/resources/requests/cpu","value":"0.5m"}]'# kubectl patch deployment -n tagify tagify --type
"json" -p '[{"op":"add","path":"/spec/template/spec/containers/0/resources/requests/cpu","value":"0.5m"}]'
```

2. Rerun upgrade.

Resolved Issues

This section lists the issues resolved in Juniper Paragon Automation Release 2.3.0:

- Sometimes, the onboarding workflow might restart for a device that is already onboarded. This is harmless, as the workflow will observe that the node is already onboarded, report the observation, and exit.
- Airflow scheduler pod may crash when multiple onboarding DAGs are triggered simultaneously.

By default, Airflow allocates 10 GB of disk space in the **PVC mount** directory (`/opt/airflow/mount/`). When multiple onboarding DAGs are triggered simultaneously, excessive logs are generated thereby filling up the disk space and potentially causing the Airflow scheduler pods to crash.

- You cannot release a Cisco device using the **Release Router** option. This issue occurs if you have added the Cisco device to Paragon Automation from the Inventory page (**Inventory > Devices > Network Inventory**) and you have used non-alphanumeric characters for the MAC address.
- For PTX10001, PRX10004, and PTX10016 devices, the PSM Temperature graph on the Hardware Details for *Device-Name* page (**Observability > Troubleshoot Devices > Device-Name > Hardware accordion > PSUs**) does not display any data.
- Sometimes, data points on the graphs on the Input/Output Traffic Details for Device-Name pages (**Observability > Health > Troubleshoot Devices > Device-Name > Overview > Interfaces (accordion) > Input/Output Traffic data link**) are displayed as critical even though the traffic is below the threshold level.

When the LACP state of a child interface is down, the aggregated interface speed also comes down based on how many child interfaces are active at the point of data collection. However, when the LACP state of the child interface is up, the aggregated interface speed also increases. Because the child interface LACP state is toggling, the aggregated interface speed also toggles and alerts are raised as the input/output traffic utilization is beyond the threshold with the reduced speed. But in the graph, since the latest thresholds are plotted, alerts might be raised even though the traffic is below the threshold.

- The Connectivity accordion (**Observability > Troubleshooting Devices > Device-Name**) shows no data although connectivity tests are properly run and passed.
- You cannot delete a Test from the Test-Name page (**Observability > Active Assurance > Test > Test-Name > *Test-Name* > More > Delete**). The following error is displayed and the Test that you tried to delete is retained:

Failed to delete Test

- On rare occasions, e-mails might not be sent out due to Kafka connection issues. The mail service logs display the following error:

"message": "dial tcp 10.x.x.43:9092: connect: connection refused"

This issue can occur anytime, but more commonly when the cluster has been upgraded or when a cluster node has been rebooted.

- The following limitations are seen when you use the service orchestration cMGD CLI to modify the placement-interface information of an L3VPN service:
 - The initial placement-interface options that were populated when the service order was created are not displayed.
 - You can select the interface for the site access from all the interfaces present on the CE or PE device.
 - When you modify the PE topology and the available ports in the topology, you must:
 1. Delete the existing placement-interface and placement-options from the site network access by using either REST API or the service orchestration cMGD CLI.
 2. Execute the request service order modify command to regenerate the service order with the modified values for the placement-options.
- While configuring an EVPN service order, the GUI does not throw a validation error even if you specify a value that is equal to 1 Tbps for CBS and CIR fields.
- Scheduling provisioning of service orders is a Beta feature in Release 2.1.0. Except in fresh installations, scheduling may not work consistently.

- While modifying a resource instance, if you update VLAN with a value higher than the current specified value then the Modify Resource Instance operation fails.
- The service order fails with the error message, Invalid XML document, namespace is missing.
- On the Customer Inventory page (**Orchestration > Customers**), when you click the **View Instances** hyperlink to view the service instances provisioned for the customer, a filter for the Customer field is automatically applied on the Service Instances page. Since filtering is not supported for the customer field, you may not see any data.
- While creating an L3VPN service instance if you click **Reset Placements**, the resources are cleared from the UI. However, these resources are not released from the back end.
- When you click **Update Placements** while modifying services, a confirmation message appears that the placements are successfully updated even though the REST API returns an error.
- You must configure access diversity while creating an EVPN service order with all-active or single-active redundancy modes. However, the Access Diversity section is not populated in the UI.
- After you upgrade Paragon Automation from Release 2.1.0 to Release 2.2.0, the *vpn* resource instance does not appear on the Resource Instances page (**Orchestration > Service > Resource Instances**). This issue occurs because the *vpn* service design has been renamed to *vpn-resources* in Release 2.2.0.
- While provisioning an EVPN service order, if you mark an access interface as untagged, the placement may still pick up a tagged interface from interface resources defined in the topology file.

For example, if et-0/0/2 is defined as untagged and et-0/0/3 is defined as a tagged interface in the topology resources, then the service order sets the access interface as untagged and the automatic placement may still pick up the tagged et-0/0/3 interface.
- If you modify an L3VPN service to update the routing protocol then the changes are not reflected in the L3VPN accordion (**Orchestration > Instances > Service Instances > *Service-Instance-Name* hyperlink > Passive Assurance**).
- We recommend that you do not enable LDP for IPv6 for the L2 circuit.
- Sometimes, the Interface Status column on the Passive Assurance tab of the *L3VPN-Name* Details page (**Orchestration > Instances > *Service-Instance-Name* hyperlink**) may not display data. Occasionally, you may not see any data on the Interface Status graph.
- Sometimes you may notice that the previously-available Test Agents and previously-provisioned Monitors are not shown in the GUI. In addition, the performance tests are timing out and the measurements are stopped.

- If you try to adopt Cisco devices using the GUI or try to register Cisco devices in bulk using the REST APIs then the system does not trigger a connection to the devices. Therefore, the device status is shown as disconnected.
- The order state of a network implementation plan does not reflect the onboarding status of individual devices in the plan. The order status may show as a success even if the device onboarding has failed.
- On the Tests page (**Observability > Active Assurance**), the Test summary that is displayed on the info card is not based on the time range that you have selected. Instead, the Test summary is based on the number of Tests listed on a specific page.