Release Notes

Published
2024-11-18

Juniper Paragon Automation

Software Highlights

- **Device Life-Cycle Management**—Automates and secures onboarding, managing, and monitoring devices.

  In this release, you can onboard Cisco Systems routers, customize device and interface provisioning during onboarding, configure device bandwidth and access parameters in the network implementation plan, and configure aggregated Ethernet interfaces.

- **Observability**—Provides actionable insights on device and network health, anomaly notifications, and an intuitive and multidimensional view of the network topology.

  In this release, you can monitor routing health, view real-time changes in topology, and discover LSPs in your network.

- **Trust and Compliance**—Monitors network targets, provides information about potential vulnerabilities, and recommends corrective actions for non-compliant devices.

  In this release, you can run a compliance scan by uploading a JSON file.

- **Service Orchestration**—Supports provisioning of network services from the Paragon Automation GUI, the service orchestration cMGD CLI, or programmatically by using REST APIs.

  In this release, you can provision EVPN and L2 circuit services. You can create and manage Resource Instances.

- **Active Assurance**—Validates the network's configuration from the end-user or from the application perspective by generating synthetic traffic from different points in the network.

  In this release, you can create Test templates, input variables, and use additional native plug-ins (UDP, TCP, IPTV MPEG, Netflix Speedtest, and OTT-HLS).

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® 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, 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. In addition, there are pre-deployed Test Agents in certain clouds, such as AWS. 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 5.

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 2.1.0.

- **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, and 23.4R2. • Junos OS releases 23.2R2, 22.4R2, 22.2R3, and 23.4R2.

Supported Juniper Devices	• ACX7024 • ACX7024-X • ACX7100-32C • ACX7100-48L • ACX7348 • ACX7509 • PTX10001-36MR • PTX10004 • PTX10008 • PTX10016 • MX204 • MX240 • MX304 • MX480 • MX960 • MX10003 • MX10004 • MX10008
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

Supported Browsers	The latest version of Google Chrome, Mozilla Firefox, and Safari.

New Features

IN THIS SECTION

This section describes the features available in Juniper Paragon Automation Release 2.1.0.

Device Life-Cycle Management

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

Juniper Paragon Automaton Release 2.1.0 provides the following additional device life-cycle management features:

- **Support for aggregated Ethernet interfaces**—Apart from channelized interfaces, management interfaces, and interfaces with logical units, you can configure aggregated Ethernet interfaces in a network implementation plan.

[See Add Network Implementation Plan.]

- **Customize device and interface provisioning during onboarding**—You can use configuration templates within device profiles and interface profiles to configure the device infrastructure during device onboarding. Additionally, you can use the templates to customize the configurations provided in the device profiles and interface profiles.

  [See Add Network Implementation Plan.]

- **Support for Cisco Systems routers**—Apart from Juniper Networks' routers, you can onboard routers from Cisco Systems to Paragon Automation and manage them. In this release, support for non-Juniper devices is limited to onboarding the device and configuring the device using REST APIs. See **Help > API Docs** in the Paragon Automation GUI for information about Paragon Automation REST APIs.

  [See "Supported Junos OS Releases, Devices, and Browsers " on page 3 for the list of devices supported in Paragon Automation.]

- **Filter devices and interfaces by using tags**—You can assign tags in the key:value format (for example, site:London) to devices and interfaces. You can use the tags to select a set of devices and interfaces based on the tags assigned to them.

  You can assign tags to devices from the Inventory page (**Inventory > Devices > Network Inventory**) and to interfaces from the Interfaces tab (**Observability > Health > Troubleshoot Devices > *Device-Name* > Inventory > Interfaces**).

  You can view, add, edit, and delete all the tags in an organization from the Tags page (**Inventory > Common Resources > Tags**).

  [See About the Tags Page.]

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 2.1.0 provides the following additional observability features:

- **Monitor routing health**—Juniper Paragon Automation monitors the overall health of routing components during device onboarding and when a device is operational. You can view the following information on the Routing accordion of the *Device-Name* (Observability > Health > Troubleshoot Devices > *Device-Name*) page:

  - Overall health status and relevant events
  - BGP peers, flaps, routes (advertised and received)
  - IS-IS adjacency state, flaps, IS-IS drops
  - OSPF interface state, hello protocol, input/output (I/O) errors
  - RSVP neighbors, global traffic engineering (TE) errors, TE interface errors
  - Label-switched path (LSP) and LDP peers, LSP flaps
  - Total routes, and total active routes in the routing information base (RIB) also known as routing table, and forwarding information base (FIB) also know as forwarding table

  [See Routing Data and Test Results.]

- **View live topology updates**—You can view live changes in topology even if the devices are not managed by Paragon Automation.

  To view live updates, you need to do the following:

  1. Specify the IP address of the BGP-LS peer and autonomous system (AS) number on the Topology Settings page (**Observability > Network > Topology > Topology Menu Bar > Settings icon**).
  2. Manually refresh the Topology page or you need to right-click on the topology map and select **Reload Network**.

  In addition, on the topology map, you can view the latest operational status of devices, links, sites, and tunnels.

  [See Network Topology Visualization Overview.]

- **Discover label-switched paths (LSPs) in your network**—Juniper Paragon Automation uses Path Computation Element Protocol (PCEP) to discover LSPs in your network. You can view all LSPs and their attributes on the Tunnels tab of the network information table on the Topology page (**Observability > Network > Topology**). You can also view the operational status of the LSP on the topology map of the Topology page.

  For tunnel-related information to be displayed on the Tunnels tab, you must adopt or onboard devices to Paragon Automation.

  [See About the Tunnels Tab.]

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 2.1.0 provides the following additional trust and compliance features:

- **Run a compliance scan by uploading a JSON file**—You can run a compliance scan by uploading a preconfigured compliance configuration file in the JSON format, in the Create Compliance Scan wizard (**Trust > Compliance > Compliance Scan > + Add > Create Compliance Scan**).

  [See Perform Custom Compliance Scan.]

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. The instruction to create, modify, or delete a service instance is a service order. After you initiate a service order and publish it, Paragon Automation provisions 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 2.1.0 provides the following additional service orchestration features:

- **Provision EVPN services**—Paragon Automation provisions Layer 2 Ethernet VPN (EVPN) services in your network by using automated workflows. To provision an EVPN service, you must define EVPN service elements such as the VPN service topology and type, and site-specific details such as site names and locations, and site network access parameters.

  You can monitor the workflow execution status and detailed task logs to troubleshoot and fix errors when an automated workflow run fails. After the EVPN service is provisioned, Paragon Automation automatically monitors the health and quality of the service.

[See Add an EVPN Service Instance.]

- **Provision L2 circuit services**—You can use Paragon Automation to provision a point-to-point L2 circuit between two customer edge devices through the MPLS network. To provision an L2 circuit, you must define L2 circuit service elements such as the underlay transport type, details about the provider edge and customer edge devices through which the L2 circuit spans, site network access parameters, and signaling type in the MPLS network.

  Paragon Automation provisions L2 circuit services by using automated workflows. You can monitor the workflow execution status and detailed task logs to troubleshoot and fix errors when a workflow run fails. You can also monitor service health and quality after the service is provisioned.

  [See Add an L2 Circuit Service Instance.]

- **Create and manage Resource Instances**—You can use Paragon Automation to create, modify, and delete resource instances based on guidelines and templates defined in the corresponding resource designs. A resource instance defines the elements of a network resource pool that must be configured to provision Layer 3 VPN (L3VPN), EVPN, and L2 circuit services in the network. After you create or modify a resource instance, and commit the instance, a service order is generated. The service order activates the automated workflow to upload the resource pool to the Paragon Automation database. A delete service order activates the automated workflow to delete the resource pool from the database. You can also monitor the execution state of the service orders and view details of the workflow run tasks associated with the service orders.

  [See About the Resource Instances Page.]

- **View service monitoring data**—Paragon Automation automatically monitors service health and quality after provisioning the service in the network. You can view the monitoring data under the Active Assurance (for L3VPN and EVPN services) and Passive Assurance (for L2 circuit services) tabs on the *Service-Instance-Name* Details page.

  [See View Service Instance Details.]

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 Junos OS Evolved routers, x86 hardware, or on virtual machines. If you are using Juniper Networks®

MX Series Universal Routers and Juniper Networks® PTX Series Routers, Paragon Automation uses real-time performance monitoring (RPM) for collecting the metric data.

Juniper Paragon Automation Release 2.1.0 provides the following additional Active Assurance features:

- **Support for Test templates**—You can create Test templates and reuse these templates while creating Tests. Test templates eliminate the need to manually configure Task parameters and metrics evaluation criteria each time you run a Test.

  You can view and manage the Test templates from the Test Templates **(Inventory > Active Assurance > Test Templates**) page.

  [See About the Test Templates Page.]

- **Support for input variables**—You can create input variables when you configure a Step for a Test or a Test template. Input variables are reusable Task parameters that ensure that the same values are consistently used across different Steps. Input variables eliminate the need to reconfigure the parameters and minimize discrepancies.

  From the Measurement Designer **(Observability > Active Assurance > Measurement Designer > + Create blank Test > Tasks)** page, you can create input variables by specifying the parameters such as label name, advanced details, and settings.

  [See Create Input Variables.]

- **Support for additional native plug-ins**—Paragon Automation supports the following native plug-ins to evaluate the quality of services in your network:

  - IPTV MPEG
  - Netflix Speedtest
  - OTT-HLS
  - TCP
  - UDP

  [See Test and Monitors Overview.]

- **Support to retain Test results for a canceled Test**—Paragon Automation provides a cancel option on the Tests (**Observability > Active Assurance > Tests**) page to stop a running test.

  When you cancel a running Test, Paragon Automation retains the data related to the canceled Test up to the point of cancellation. You can view this data on the Tests page.

  [See About the Tests Page.]

- **View Test Agents and Monitors details**—You can view the details of Test Agents and Monitors such as name, descriptions, and tags. To view these details, a **Detailed View** icon is added on the Test Agents **(Inventory > Active Assurance > Test Agents)** page and the Monitors **(Observability > Active Assurance > Monitors)** page. In addition, you can also edit the **Name**, **Description**, and **Tags** fields, and copy the API Request URL to fetch the details of the Test Agents and the Monitors.

  [See About the Test Agents Page and About the Monitors Page.]

Administration

Paragon Automation Release 2.1.0 provides the following administration features to manage users, sites, and organizations:

- **Audit log enhancements**—Audit logs are available for user account lockout, unsuccessful authentication, and successful authentication events.

  The type of event for which the log is generated is displayed in the Event Type column of the Audit Logs page (**Settings menu > Audit Logs)**.

  Event Types include:

  - User Management (events such as user creation and deletion, password reset, successful, unsuccessful authentication, and so on).
  - Organization (events at the organization level such as creating or deleting an organization, user logging into or exiting the organization, and so on).

  [See Audit Logs Overview and About the Audit Logs Page.]

Paragon Shell CLI

The following command has been changed in Juniper Paragon Automation Release 2.1.0:

- **request paragon support information**—You can use the `request paragon support information` command to view an in-depth status report of the Paragon Automation cluster configuration. When you execute this command on any one of the cluster nodes, a series of `show` commands and `kubectl` commands are executed one after the other.

  The `request support information` command is no longer valid.

  [See request paragon support information.]

Beta Features

Juniper Paragon Automation Release 2.1.0 provides Beta support for the following features:

- **Observability**
  - **Health Dashboard**—Paragon Automation provides a dashboard that enables you to monitor network health in real-time. On the WAN Health tab of the dashboard (**Observability > Health > Health Dashboard > WAN Health**), you can view the overall health of your devices, interfaces, and routing neighbors. You can also view KPIs that affect overall health, and a graph of the average health of devices for the past 30 minutes.

    [See About the Health Dashboard and About the WAN Health Tab.]
- **Service Orchestration**
  - **Schedule service order provisioning**—You can schedule provisioning of an L3VPN, EVPN, and L2 circuit service order by specifying the date and time for provisioning. Paragon Automation automatically provisions the service order at the specified date and time.

    [See Add an L3VPN Service Instance, Add an EVPN Service Instance, and Add an L2 Circuit Service Instance.]
- **Administration**
  - **Support for Lightweight Directory Access Protocol (LDAP)**—Apart from Security Assertion Markup Language (SAML), you can configure LDAP to authenticate and authorize users logging into Paragon Automation.

    [See Manage Identity Providers.]

Deprecated Features

The following feature is deprecated in Juniper Paragon Automation Release 2.1.0.

- We do not support using the cluster deployment wizard to deploy your Paragon Automation cluster. You must use Paragon Shell to deploy your cluster.

Known Issues

IN THIS SECTION

This section lists the known issues in Juniper Paragon Automation.

Device Life-Cycle Management

- When you use the **Release Router** option to release a device from being managed by Paragon Automation, the device might not be released as the service orchestration engine might still be referencing the device that you want to release.

  Workaround: Before you use the **Release Router** option, update the network implementation plan and service instances so that the services no longer use the device you are trying to release. Also, the device should be removed from the resource pool that the services are accessing..

- 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 so that the onboarding workflow restarts.

- 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.

  Workaround: None.

- The onboarding of a device fails if you use NETCONF EDIT or NETCONF RPC configuration formats in the configuration templates.

  Workaround: NETCONF EDIT and NETCONF RPC configuration formats are not supported in Release 2.1.0. Instead, use the CLI configuration format.

- Paragon Automation triggers the configuration templates included in a device profiles 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: None.

- On the Inventory page (**Inventory > Devices > Network Inventory**) page, the device status is displayed as `Connected` even if there is a network disconnection in outbound SSH. Ideally, the device status should be displayed as `Disconnected` within 10 minutes. Sometimes, there may be a delay for the changes to be reflected on the Inventory page.

  Workaround: None.

- A warning message is not displayed when you try to delete a configuration template that is used in the network implementation plan.

  Workaround: None.

- 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 **Administration > Settings** in Paragon Automation 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
```

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.
  - Line card charts for some ACX Series and MX 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
MX Series	MX204, MX240, MX304, MX480, MX960, MX10004, MX10008, MX10016	fpc-temperature, fpc-cpu-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 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.

- For PTX10001, PRX10004, and PTX10016 devices, the Linecards graph on Hardware details for *Device-Name* (**Observability > Troubleshoot Devices > Device-Name > Hardware accordion**) page does not display any data.

  Workaround: None.

- For PTX Series, MX Series, and ACX Series devices, the RSVP TE Global Errors graph on the RSVP Routing Details for *Device-Name* (**Observability > Troubleshoot Devices > Device-Name > Routing and MPLS accordion**) page does not display any data.

  Workaround: None.

- 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.

  Workaround: None.

- 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.

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.

- 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.

- 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/

- If you modify an L3VPN service, the historical data for Monitors related to the modified L3VPN service are deleted.

  Workaround: None.

- The history of service orders that are generated for a service instance is saved in the Order History Tab for auditing purposes. However, when you delete the service instance, the service order history gets deleted.

Workaround: None.

- There is no synchronization between accessing and configuring devices. A workflow might fail if a device is configured by more than one service order.

  Workaround: If two different service orders are likely to affect the same device, we recommend that you wait until the first service order is executed before you publish the next service order.

- 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.

  Workaround: Based on your topology, ensure that you specify the right values for CBS and CIR fields.

- 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.

- Sometimes, the publishing of a service order fails due to existing placements.

  Workaround: In such cases, 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, discard the placements and then 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.

- When you edit a topology resource instance, the POP page may not list the latest sites or nodes.

  Workaround: Refresh the Resource Instances page before you start the edit operation.

- The publishing of an EVPN service order fails if you modify an existing EVPN service order to add a new site. This issue occurs only when you modify an existing EVPN using the GUI.

  Workaround: If you want to add a new site to an existing EVPN service order:

  1. Export the service order and save the service order in JSON format on your local system.
  2. Edit the service order and import the service order from your local system.
  3. Save and publish the service order.

- All access circuits must have the same VLAN configured, failing which the service may not function as desired.

  Workaround: None.

- Scheduling provisioning of service orders is a Beta feature in Release 2.1.0. Except in fresh installations, scheduling may not work consistently.

  Workaround: None.

- The **VPN Node Id** field on the VPN Node page (**Service Orchestration > Instances > Add L2 Circuit**) field may not list all the onboarded devices.

  Workaround: Refresh the Instances page before you add the L2 circuit service instance.

- Even though the EVPN service order is successfully created and the configurations are pushed to the devices, the following error message is displayed and the communication between the customer Edge (CE) devices fails:

  ```
  Status: Configuration Error
  ```

This issue occurs in a multi-homing scenario (multiple provider edge (PE) devices are connected to the same customer edge device) with *single-active* or *all-active redundancy* mode.

Workaround: Configure a unique LAG index for each site.

- 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.

- While creating an EVPN site, if the value that you have specified for the **Minimum number of links** field is greater than the number of Member Links, then the EVPN service order fails.

  Workaround: Ensure that the value you specify for the **Minimum number of links** field is less than or equal to the number of Member Links configured for the LAG interface.

- 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.

  Workaround: None.

- If you try to modify an existing resource instance that does not have a device or a site, then the Inventory section of the Modify *Resource-Instance-Name* page does not load.

  Workaround: Ensure that you have added at least one device and a site to the resource instance.

- The service order fails with the error message, `Invalid XML document, namespace is missing`.

  Workaround: On the device with the failed configuration, you should turn off `system services netconf rfc-compliant` and `system services netconf notification` knobs.

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

  Workaround: Instead of Microsoft Azure and Google Cloud Platform, use Amazon Web Services (AWS) as the endpoint.

- 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 the action to be taken as `Drop` when the upper limit for customer MAC addresses exceed.

  Workaround: None.

- When you modify a resource instance, Link Aggregation Control Protocol (LACP)-related information is not displayed.

  Workaround: Use the REST API to upload the LACP-related resource. See **Help > API Docs** in the Paragon Automation GUI for information about Paragon Automation REST APIs.

Active Assurance

- The REST API, **api-aggregator**, does not capture alerts that are related to Tests on the Connectivity accordion (**Observability > Troubleshooting Devices > Device-Name**).

  Workaround: None.

- 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.

  Workaround: None.

- If you modify an existing Monitor and then restart the Monitor, the events that are raised before modifying the Monitor are not cleared.

  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.

  Workaround: Reinstall Test Agent after the Routing Engine switchover.

Administration

- The latest audit log messages may not be displayed on the Audit Logs page.

  Workaround: Restart the audits-delivery stream deployment. To restart the audits-delivery stream deployment:

  1. Log in to a Paragon Automation cluster node.
  2. Run the following commands:

     ```
     kubectl -n streams scale --replicas=0 deployment audits-delivery
     ```

```
kubectl -n streams scale --replicas=1 deployment audits-delivery
```

3. Check whether the latest audit logs are listed on the Audit Logs page. If you don't see the latest audit logs, you may have to repeat this procedure.

Installation

- The backup and restore functionality has the following caveats:
  - You cannot restore backed up data from a Release 2.0.0 setup to a Release 2.1.0 setup.
  - You must restore data on the same setup from which the backup was taken and not on a fresh installation.
  - 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 the PCE Server VIP address is not configured, `kube-proxy` is set to a random port.

  Workaround: Configure the PCE Server VIP address.
- 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 three, 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* state and are not in any intermediate states like *Terminating*, *CrashloopbackOff*, and so on.

- If there are multiple node failures, the OpenSearch database may fail to start and rejoin the cluster.

  Workaround: Manually restart the OpenSearch database by running the `kubectl rollout restart sts -n common opensearch-cluster-master` command.

  After all three OpenSearch instances are restarted, monitor the log of each OpenSearch instance to see if the pod logs do not have any obvious error message. In rare cases, you may need to restart OpenSearch multiple times.

- 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
  ```

Resolved Issues

This section lists the issues resolved in Juniper Paragon Automation Release 2.1.0:

- You cannot delete a Monitor from the *Monitor-Name* page (**Observability > Active Assurance > Monitors > *Monitor-Name***).

- When you upload the topology network resources for a service, you must ensure that the interfaces are not channelized; otherwise, the placement fails.

- The following ports are reachable from outside the Paragon Automation cluster; we recommend that you block access to these ports: 22, 443, 2222, 2379, 2380, 5473, 6443, 7472, 7946, 8443, 9345, 10250, 10260.

  NOTE: These ports should not be blocked for intra-cluster communication.

- The `request service project add /data-models/projects/project-name` command fails if the project that is being added has a dependency on another project. For example, if the **onboard.tgz** has a dependency

on **network-resource.tgz**, then the `request service project add /data-models/projects/onboard.tgz` command fails.

- When you enable the **Show LEDs, Ports & Cables on Chassis** toggle button on the Hardware accordion (**Observability > Troubleshoot Device >** ***Device-name*** **> Overview** Tab), the chassis view displays the port status incorrectly.
- On the Hardware accordion (**Observability > Troubleshoot Device >** ***Device-Name***), the value of available units for CPU and Memory is always displayed one extra compared to the units that are actually available on the device. This issue occurs for the following devices: ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7509, ACX7348, MX204, MX240, MX304, MX480, MX960, MX10004, MX10008, MX10016.
- While configuring a Site Network Access for an L3VPN service, the minimum value that you must specify for **Service Input Bandwidth** and **Service Output Bandwidth** fields must be 20,000 bps. Otherwise, the service provisioning fails.
- The customer name and the VPN ID must be unique when you create an L3VPN service order. Otherwise, the service order creation fails.
- If an adopted device and a node that is not yet onboarded have the same serial number, then they are listed as two different devices on the Put Device Into Service page (**Inventory > Onboarding Dashboard**).
- The Optical Rx Power and Optical Tx Power charts do not display any data when you click the Input Traffic data-link on the Interfaces accordion (**Observability > Health > Troubleshoot Devices >** ***Device-Name*** **> Overview**).
- While creating or modifying an L3VPN service, you cannot set the Access Diversity settings parameters on the Site Network Access page and set the device references simultaneously. The placement fails to place the site network access on an available port.
- When you navigate to chassis view (**Observability > Troubleshoot Device >** ***Device-name*** **> Overview Tab > Hardware** accordion) and either try to access the minimize/maximize controls or wait for a minute or two, sometimes an “Error occurred in the component” message is displayed and the entire Overview tab blanks out.
- The VPN ID field on the Add or Modify L3VPN page is restricted to 9 characters to avoid failure of L3VPN service provisioning. This restriction is due to a configuration field size limit on Junos OS.
- When you run the `request paragon ssh-key` command, the previously-generated SSH keys in **.ssh/authorized_keys** and the SSH keys that you have added are cleared, and new SSH keys are generated for nodes.
- Paragon Automation does not push firewall filter configuration to ACX Series devices as firewall filter configuration is not supported on them.

- If you modify an existing Monitor and then restart the Monitor, the events that are raised before modifying the Monitor are not cleared.
- If you make multiple changes to a service order before executing a service order then only the last service order has a workflow associated with it. When you delete a service instance the order history also gets deleted.
- The output of the `request paragon storage cleanup` command erroneously displays the total reclaimed space as 0 bytes even though multiple files are cleared after executing this command.
- If you modify an existing L3VPN service to update certain fields such as access diversity, zip code, or device reference, the service provisioning fails as the outdated placement-related information is sent to the REST API.
- When SMTP is configured through Paragon Shell, and you enable e-mail notification for alerts and alarms, the e-mails are sent from *no-reply@mailservices.juniper.net*, instead of what is configured through Paragon shell.
- When you enable e-mail notification for alerts and alarms, e-mails are sent to the recipients even if system-wide SMTP is not configured through Paragon Shell.
- You might be able to execute REST APIs that are not supported.
- Initially, no data is displayed on the Fans charts when you select time ranges other than **30m** on the Hardware Details for *Device-Name* page in the Hardware accordion (**Observability > Troubleshooting Devices > Device-Name**). Based on the time range you have selected, you need to wait for the following time duration for the data to be populated:
  - **3h**—Approximately 10 minutes.
  - **1d**—Approximately 1 hour.
  - **1w**—Approximately 7 hours.
- Even though the Delete L3VPN Service operation is successful, the service order and the resources allocated to the service are not removed.
- When you modify the VPN ID for an existing L3VPN service, the provisioning fails because the PAA RPM monitor refers to the old routing instance.
- The service orchestration engine onboards nodes in a network implementation plan sequentially even if the devices are adopted simultaneously. The service orchestration engine can onboard devices in different plans concurrently. We recommend that you do not onboard more than 3 devices, spread across 3 different plans, concurrently.