

Release Notes

Published
2023-04-18

Paragon Automation, Release 21.3

Table of Contents

Introduction to Paragon Automation	1
Installation and Upgrade Instructions	1
Licensing	2
New and Changed Features	2
Known Issues	6
Resolved Issues	19

Introduction to Paragon Automation

Juniper® Paragon Automation is a cloud-ready solution for network planning, configuration, provisioning, traffic engineering, monitoring, and life-cycle management that brings advanced visualization capabilities and analytics to network management and monitoring. You can deploy Paragon Automation as an on-premises (customer-managed) application.

Paragon Automation operates on a microservices-based architecture and employs REST APIs, gRPC APIs, and common messaging bus communications. Paragon Automation provides base platform capabilities such as support for Juniper Networks and third-party (Cisco IOS XR, Nokia) devices, zero-touch provisioning, user management and role-based access control (RBAC), and so on. In addition to providing base platform capabilities, Paragon Automation offers users a suite of microservices-based applications—Juniper® Paragon Insights (formerly HealthBot), Juniper® Paragon Planner (formerly NorthStar Planner), and Juniper® Paragon Pathfinder (formerly NorthStar Controller). When you add any of these applications to Paragon Automation, the API suite of the application integrates with Paragon Automation to allow seamless communication between new and existing services.

The solution is an open architecture that allows integration with third-party software. Paragon Automation supports out-of-the-box integration for Juniper Networks' partner application, Anuta ATOM, which provides advanced workflow management and the service provisioning capabilities.

In these release notes, we outline the features of the base platform, Paragon Pathfinder, Paragon Planner (Web Planner and Desktop Application), and Paragon Insights modules that are available in this release. For more information about features related to these applications, see [Paragon Automation User Guide](#).

Installation and Upgrade Instructions

For information about installation procedure, upgrade procedure, and requirements (software and hardware), see [Paragon Automation Installation Guide](#).

NOTE: If your installed version is Paragon Automation Release 21.1, you cannot upgrade to Paragon Automation Release 21.3. You must perform a fresh installation of Paragon Automation Release 21.3.

If your installed version is Paragon Automation Release 21.2, you can upgrade to Paragon Automation Release 21.3.

Licensing

In Paragon Insights, the following license tiers and their related device licenses are enforced:

- Paragon Insights Advanced (PIN-Advanced)
- Paragon Insights Standard (PIN-Standard)

Currently, the tier licenses are hard enforced. That is, you cannot perform the deploy operation unless you add the licenses. The device licenses are soft enforced. That is, you will receive an out-of-compliance alert in the Paragon Automation GUI if you exceed the number of devices for which you have obtained licenses. However, the existing functionality will not be blocked. You can view your license compliance status on the **Administration > License Management** page in the GUI.

In Paragon Pathfinder, the following license tiers are enforced:

- Pathfinder Standard
- Pathfinder Advanced
- Pathfinder Premium

For information about licensing, see the [Licensing Guide](#).

Starting in Paragon Automation Release 21.3, the Paragon Automation icon in the top-left corner of the GUI is updated depending on the license that you add. When you log in to the Paragon Automation GUI for the first time, the Paragon Automation icon appears without the names of the three components below it. The GUI displays the name of a component only after you add a license for that component. For example, after you add a Paragon Insights license, the name Insights appears below the Paragon Automation icon. After you add a license for Paragon Pathfinder, the names Pathfinder and Planner are also displayed.

New and Changed Features

This section describes the features in each module of Juniper Paragon Automation Release 21.3.

Paragon Insights

- **Enhancements to committing configuration changes**—In Paragon Automation Release 21.3, when you commit a configuration, the changes are immediately accepted after validation. A background job that runs periodically tracks these changes and applies the changes to ingest services. Only after the background job applies these changes, does the ingest service process data based on the changes

configured. In releases prior to Paragon Automation Release 21.3, commit configuration changes are accepted and acknowledged only after the changes are applied to ingest services.

[See [Commit or Roll Back Configuration Changes in Paragon Insights](#).]

- **Manage action engine workflow commands, conditions, inputs, and outputs**—Action engine workflow continues to be a Beta feature in Release 21.3. Starting with Paragon Automation Release 21.3, you can add or edit commands, conditions, inputs, and outputs in an action engine workflow from the **Configuration > Action Engine** page.
[See [Manage Action Workflows](#).]
- **Monitoring clusters using kube-state-metrics**—kube-state-metrics service is a beta feature in Paragon Automation Release 21.3. kube-state-metrics is a third-party service that generates metrics based on the current state of Kubernetes clusters. kube-state-metrics runs as a cluster service, and is automatically installed when you install Paragon Automation.
[See [Understand Kube-State-Metrics Service](#).]
- **Support for unsigned integer data type**—Starting in Paragon Automation Release 21.3, you can select unsigned integer as a data type when you add a Paragon Insights rule, a raw data summarization profile, or a tagging profile. An unsigned integer is a data type that can contain values from 0 through 4,294,967,295.
[See [Configure a Custom Rule](#), [Add a Raw Data Summarization Profile](#), and [Add a Tagging Profile](#).]
- **Linking action engine workflows to rules**—Action engine workflow continues to be a beta feature in Release 21.3. Starting in Paragon Automation Release 21.3, you can link action engine workflows to rules when you set up triggers. However, you need a PIN-Advanced license to use this feature.
[See [Configure a Custom Rule](#).]
- **Support for third-party vendor devices when you configure rule properties**—Starting in Paragon Automation Release 21.3, you can add vendor identifier, vendor name, default sensors, product, platform, release, and operating system-related information for third-party devices when you configure rule properties.
[See [Configure a Custom Rule](#).]
- **Applying default sensors to devices**—Starting in Paragon Automation Release 21.3, when you configure a Paragon Insights rule, you can select the default sensors to be applied to all supported devices, only to Juniper devices, only to Juniper devices that run a specific OS, or to all third-party (non-Juniper) devices.
[See [Configure a Custom Rule](#).]
- **Enhancements to Grafana UI for better usability**—Starting in Paragon Automation Release 21.3, you can apply Paragon Insights filters to view aggregate data of a device group, or view aggregate data of multiple device groups from the Grafana UI. The Juniper Paragon Insights TSDB plug-in that is available with Paragon Automation Release 21.3 enables you to apply these filters from the Grafana UI.

[See [Grafana Overview](#).]

- **View and configure resources and dependencies in the network**—Starting with Paragon Automation 21.3, you can get view the resources and dependencies between the resources in your network. Resources can be any device components (for example, chassis or line card) or network components that span many devices (for example, VPN and IPSec). Dependencies between resources show how events in one resource impact other resources, enabling you to analyze the root cause of failures triggered by multiple events.

[See [Understand Resources and Dependencies](#).]

- **Support for smart alarms**—Starting with Paragon Automation Release 21.3, you can view smart alarms on the Alerts page, generated by resource dependencies. Smart alarms is a feature where alarms from different rules are displayed in a tree structure, where the top-level alarm represents the root cause of the other alarm events. This feature was also available in the Beta release of Paragon Automation Release 21.3.

[See [About the Alerts Page](#).]

- **Support for Bring Your Own Ingest (BYOI) default plug-ins**—Starting with Paragon Automation Release 21.3, you can work with Juniper Networks to design default BYOI plug-ins. A default BYOI plug-in processes telemetry data collected by other collectors and sends the data to the Paragon Automation pipeline for analysis. The BYOI plug-in ingest can import data from different sources such as data lakes that have different data models, different data encoding, and security.

[See [Understand Bring Your Own Ingest](#).]

- **Support for multiple return values in user-defined functions**—Starting with Paragon Automation Release 21.3, you can add multiple return values in user-defined functions (UDF). You can configure fields to capture extra return values in the Functions tab of the Rules page. The first return value is stored in the rule field. You can use this feature to reduce the number of UDF files you need to upload in Paragon Automation.

[See [Functions](#).]

- **Support for splat operator in ICMP outlier detection playbook**—In outlier detection playbook, users used to provide device IDs of each device in a device group in the XML Path Language (XPath) fields. Starting with Paragon Automation Release 21.3, you can use the splat operator instead of individual device IDs in the round-trip time (RTT) XPath field of the ICMP outlier detection playbook instances.

[See [Create a Playbook Using the Paragon Insights GUI](#).]

- **Support for enabling and disabling the Port field on the Ingest Settings page**—In Paragon Automation 21.3, you can enable or disable the outbound SSH port for iAgent or the port for native Google Protocol Buffer connections by using the toggle button on the **SSH** and **Native GPB** Ingest Settings tabs of the **Configuration > Data Ingest > Settings** page.

If you enable the Port field, you must specify the port number for SSH or native GPB connections. If you disable the Port field, the port number is cleared.

[See [Configure Native GPB Ingest](#)[Configure Outbound SSH Port for iAgent.](#)]

Paragon Pathfinder

- **Additional ingest methods to collect analytics data**—Paragon Automation Release 21.3 supports the following additional ingest types to collect analytics data:
 - SNMP to collect data from Cisco IOS XR devices and Nokia devices.
 - NetFlow to collect data from Juniper Networks devices and Cisco IOS XR devices.

[See [Collect Analytics Data Overview.](#)]

- **Support for the NetFlow collector**—Starting in Paragon Automation Release 21.3, Paragon Pathfinder supports the NetFlow collector. You can use this data collection tool to monitor the traffic flow and generate demand reports that provide information about network traffic.

[See [NetFlow Collector Overview.](#)]

- **Support for routing LSPs within network slices by using policy-based path computation profiles (Beta)**—Starting in Paragon Automation Release 21.3, you can use policy-based path computation profiles to route label-switched paths (LSPs) within a network slice.

[See [Configure LSP Routing by Using a Computation Profile.](#)]

- **Add a Test Agent for network slices**—Starting in Paragon Automation Release 21.3 you can use Paragon Pathfinder to visualize connections between Paragon Active Assurance Test Agents and devices in a network slice. The Test Agents generate active, synthetic traffic to measure service quality in the network slice.

[See [Add a Test Agent for Network Slices.](#)]

Base Platform

- **Export inventory details of a device**—Starting in Paragon Automation Release 21.3, you can export a device's inventory details to a ZIP (.zip) file. The ZIP file contains these comma-separated value (CSV) files: **chassis.csv**, **feature.csv**, **interface.csv**, **license.csv**, and **software.csv**.

If any of the inventory resources data is empty, the corresponding CSV file is not included in the ZIP file.

[See [About the Devices Page View Device Inventory.](#)]

- **Support for Nokia devices**—Starting in Paragon Automation Release 21.3, you can onboard and manage Nokia 7250 and Nokia 7750 devices that are running TiMOS version 19.10 or later.

[See [Devices Overview.](#)]

Paragon Planner

- **Support for Web Planner (Beta)**—Starting in Paragon Automation Release 21.3, the Paragon Planner Web Planner application is supported. You can launch the Paragon Planner Web application from

Planning > Paragon Planner. With Web Planner, you can visualize your network elements, forecast the impact of network changes (such as latency, traffic shifts, and new capacity) on transport services, or run failure simulations, without affecting the live network.

[See [Paragon Planner Overview](#) in *Paragon Automation Web Application User Guide*.]

Graphical User Interface (GUI)

- **Paragon Automation icon to reflect the licenses installed**—Starting in Paragon Automation Release 21.3, the Paragon Automation icon in the top-left corner of the GUI is updated depending on the license that you add. When you log in to the Paragon Automation GUI for the first time, the **Paragon Automation** icon appears without the names of the three components below it. The GUI displays the name of a component only after you add a license for that component. For example, when you add a Paragon Insights license, the name **Insights** appears below the Paragon Automation icon. When you add a license for Paragon Pathfinder, the names **Pathfinder** and **Planner** are also displayed.

[See [Paragon Insights Licensing Overview](#) .]

Paragon Installation and Upgrade

- **Support for upgrading Paragon Automation to Release 21.3**—You can upgrade your Paragon Automation Release 21.2 cluster to Release 21.3. However, you cannot upgrade from Release 21.1 to Release 21.3.

[See [Upgrade to Paragon Automation Release 21.3](#).]

- **Support for customizing the inventory file**—In Paragon Automation Release 21.3, you can customize the inventory file by using the `./ run config-dir inv` command. The `inv` command launches a configuration wizard that lets you enter cluster information into the **inventory** file. In earlier releases, you could edit the **inventory** file only using a Linux text editor.

Anuta ATOM

- There are no new features related to Anuta ATOM in Paragon Automation Release 21.3.

Known Issues

IN THIS SECTION

- [Installation | 7](#)
- [General | 8](#)

This section lists the known issues in Juniper Paragon Automation Release 21.3.

Installation

- In the absence of a time series database (TSDB) HA replication, if a Kubernetes worker node running a TSDB pod goes down, even though there is capacity in the pod, the TSDB service is not spun up on a new node. This is because a huge volume of data would need to be transferred to the new node.

Workaround: In the event of a failure of the server or storage hosting a TSDB instance, you can rebuild the server or damaged component.

If the replication factor is set to 1, then the TSDB data for that instance is lost. In that case, you need to remove the failed TSDB node from Paragon Automation. To remove the failed TSDB Node:

1. In the Paragon Automation GUI, select **Configuration > Insights Settings**.

The Insights Settings page appears.

2. Click the **TSDB** tab to view the TSDB Settings tabbed page.
3. To delete the failed node, on the TSDB Settings tabbed page, click **X** next to the name of the failed TSDB node.

NOTE: We recommend that you delete TSDB nodes during a maintenance window since some services will be restarted and the Paragon Automation GUI will be unresponsive while the TSDB work is performed.

4. Click **Save and Deploy**.
 5. If the changes are not deployed and if you encounter an error while deploying, enable the **Force** toggle button and commit the changes by clicking **Save and Deploy**. By doing so, the system ignores the error encountered while adjusting the TSDB settings.
- After you install, any changes to the **northstar.cfg** file (bootstrap configuration) are not picked up by the components.

Workaround: When you change the bootstrap configuration, use the installer to uninstall and to re-install the Paragon Pathfinder application by running the following commands:

```
./run -c <config-dir> destroy -t northstar
```

```
./run -c <config-dir> deploy -t northstar
```

- You cannot log in to the Kibana application even after successfully installing Paragon Automation because Release 21.2 supports the Open Distro version of Elasticsearch. Open Distro provides authenticated access to Kibana.

Workaround: To log in to Kibana, you **must** configure a password in the **opendistro_es_admin_password** field in the **config.yml** file before installation. For more information, see [Install Paragon Automation](#).

- If you uninstall Paragon Automation completely, you must also ensure that the **/var/lib/rook** directory is removed on all nodes, and all Ceph block devices are wiped.

Workaround: See the [Troubleshooting Ceph and Rook > Repair a Failed Disk](#) section in the Paragon Automation Installation Guide.

- At the time of installation, debug user accounts (For example, hb-northstar-admin, hb-tm-admin, hb-ems-admin) are created by default. In some deployments, the debug user accounts are not created at the time of installation, and therefore you cannot log in to the Paragon Automation GUI using the debug user account.

Workaround: You need to manually create debug user accounts. To manually create debug user accounts:

1. Get the pod ID of config-server in the healthbot namespace.

```
kubect1 exec -it $(kubect1 get pods -n healthbot -o=name | grep config-server | sed "s/^.{4\}//" ) -n healthbot bash
```

2. Log in to the config-server pod, and run the following commands to create debug users:

```
./var/local/healthbot/healthbot.py initialize-tenant -t hb-northstar
```

```
./var/local/healthbot/healthbot.py initialize-tenant -t hb-tm
```

```
./var/local/healthbot/healthbot.py initialize-tenant -t hb-ems-dmon
```

General

- Paragon Insights maps the **Name** (hostname or IP address) field to the **Device ID** field. However, the device name is no longer unique for the following reasons:
 - In a dual Routing Engine device, “-reX” is appended to the device name.
 - Third-party applications like Anuta Atom append the domain name to the device name.

Also, mapping a device by its universal unique identifier (UUID) and not the hostname could cause issues with the information that the GUI displays.

Workaround: Configure an additional IP address for the management Ethernet interface by including the `master-only` statement at the `[edit groups]` hierarchy level. You must then use this additional IP address for onboarding the device. For more information, see [Management Ethernet Interfaces](#).

- If you have dedicated a node for TSDB, some services (for example, AtomDB, ZooKeeper, and so on) in the common namespace that have `PersistentVolumeClaim` set can be affected if the relevant pods are running on the dedicated node. That is, the status of pods running on the TSDB node is always displayed as Pending.

Workaround: To avoid this situation, while dedicating a node for TSDB, ensure that the node does not have any pods for dedicated services that use `PersistentVolumeClaim`.

- The Edit Device Group operation fails if you try to add an unmanaged device to an existing device group.

Workaround: There is no known workaround.

- If you select a saved query on the Alarms page, the alarms are filtered based on the saved query. But, the graph and the date are not updated.

Workaround: There is no known workaround.

- The restore configuration operation fails for devices running Cisco IOS XR Release 7.1.1.

Workaround: There is no known workaround.

- On the Add Filter page, the Save option is not displayed if you create more than one filter on the GUI pages.

Workaround: There is no known workaround.

- While you are adding or editing a point-to-multipoint (P2MP) group, the value is not auto-populated for both MVPN Instance and Route Distinguisher fields. The values are auto-populated only for either of the fields.

Workaround: There is no known workaround.

- If you add an unmanaged device on the Device page and later edit the hostname of the unmanaged device, the hostname is not reflected in the device group and in the Devices dashlet on the Dashboard.

Workaround: You can add an unmanaged device using the hostname or the IP address of a device.

If you have added an unmanaged device using the hostname, then deleting the existing device and adding the device with a new hostname resolves the issue.

If you have added an unmanaged device using the IP address, then in the device group and the Devices dashlet on the Dashboard, you need to identify the unmanaged devices that are edited based on the IP address and not the hostname.

- While adding a filter to a table in GUI pages that support filtering, if you use both AND and OR logical operators as filter conditions, the results are not as expected.

Workaround: We recommend that you use only one logical operator for filter conditions.

- Message Digest Algorithm 5 (MD5) authentication is not supported on a Path Computation Element Protocol (PCEP) server.

Workaround: There is no known workaround.

- By default, the topology filter is disabled. You cannot enable the topology filter by using the Paragon Automation GUI.

Workaround: You can enable the topology filter by using the following procedure:

1. Log in to the ns-toposerver pod:

```
kubectl exec -ti -n northstar ns-toposerver-Pod ID-c ns-toposerver -- bash
```

2. Update the **northstar.cfg** file that is available at the **/opt/northstar/data/** location.

```
sed -i "s|^bmp_host=.*|bmp_host=ns-filter|;s|^bmp_port=10002|bmp_port=10004|;" /opt/northstar/data/northstar.cfg
```

3. Apply changes to the configMap file.

```
sed -i "s|^bmp_host=.*|bmp_host=ns-filter|;s|^bmp_port=10002|bmp_port=10004|;" /opt/northstar/data/northstar.cfg
```

4. Verify whether the Topology Filter field is enabled in the GUI.

- On the Devices page, there is no correlation between the Management Status column and the Sync Status column. For example, even if the device discovery fails, the Sync Status column might display the status as In Sync, which is incorrect. The In Sync state only represents that inventory information stored in Paragon Automation is synchronized with the device in the network.

Workaround: There is no known workaround.

- For Cisco IOS XR devices, you must set the default NETCONF port to 22, otherwise, you cannot view alarms raised on the Alarms page.

Workaround: Manually set the NETCONF port to 22.

- While adding a device, an error message is not displayed if the add device operation fails.

Workaround: There is no known workaround.

- P2MP groups configured by PCEP with flowspec mapping to multicast VPN service is not supported

Workaround: There is no known workaround.

- For Cisco IOS XR devices, you cannot restore a device configuration from the **Devices** page. You can only back up the device configuration.

Workaround: To restore the device configuration of your Cisco IOS XR devices:

1. On the **Configuration > Devices** page, select the Cisco XR device and click **More > Configuration Version**.
2. Copy the configuration version that you want to restore.
3. Restore the configuration using the CLI.

- Cisco Model Driven Telemetry (MDT) is not supported.

Workaround: There is no known workaround.

- If you perform the Undelegate operation on a delegated LSP, the Path Computation Server (PCS) uses the bandwidth reported by the device for the planned bandwidth instead of the user input value.

Workaround: There is no known workaround.

- On the **Topology > Interface** page, you cannot filter interfaces on a specific node.

Workaround: There is no known workaround.

- While you are adding a device to a device group, the Devices field on the Add Device Group page displays the UUID of the device instead of the hostname.

Workaround: There is no known workaround.

- While upgrading the image of a device, you cannot copy the image on to the device if the bandwidth on the device is lesser than 600Kbps.

Workaround: There is no known workaround.

- There is no GUI support for outbound SSH at the device level. By default, the outbound SSH is enabled at the device level.

Workaround: You can enable or disable the outbound SSH on the device by using the MGD CLI or Rest APIs. To disable the outbound SSH you must set the disable flag to true. Run the following command on the device to disable the outbound SSH using the MGD CLI:

```
set healthbot DeviceName outbound-ssh disable true
```

- You cannot download all service logs from the Paragon Automation GUI.

Workaround: You can view all service logs in Elastic Search Database (ESDB) and Kibana. To log in to Kibana or ESDB, you must configure a password in the **opendistro_es_admin_password** field in the **config.yml** file before installation.

- When you click a device on the Devices page, you cannot view the chassis details if the menu bar, which is available on the left-side of the Paragon Automation GUI, is expanded.

Workaround: Minimize the menu bar to view the chassis details.

- While adding a task for a workflow, the value that you specify for the **Every** field in the Recurrence Option section of the Container Normalization Task tab must not be lesser than 360 minutes.

Workaround: There is no known workaround.

- In Paragon Insights, if you increase the number of devices, the time taken to execute the **api/v2/config/configuration** API increases exponentially. This might impact cosmetic and administrative base platform functionalities

Workaround: There is no known workaround.

- In the **Configuration > Action Engine > Add New Workflow** page, the list of tasks in the **Entry Task** and **Exit Task** drop-down lists in the **General** tab, are not getting updated to reflect the current tasks that are created or deleted in the **Tasks** tab.

Workaround: There is no known workaround

- If you modify an existing LSP or use a slice ID as one of the routing criteria, then the path preview might not appear correctly.

Workaround: There is no known workaround

- On the Health Reports page, a blank page is displayed if you select two reports and click **Diff Reports**.

Workaround: There is no known workaround.

- The PDF report that you receive through e-mail does not include any data.

Workaround: Generate the report in HTML format. To generate reports in HTML format, on the **Configuration > Insights Settings > Add a Report Settings** page, set the **Report Format** field to HTML.

- Interfaces are deleted from the Redis database after you run the `sync topology` command.

Workaround: Rerun the device collection task.

- If a device is running a Junos OS Evolved Release, the image of the device is not upgraded during the ZTP process.

Workaround: You must ensure the configuration file is present at the time of ZTP.

- On the Add Device Group page, all devices might not be displayed in the Devices field.

Workaround: Restart config-server using the following command:

```
/var/local/healthbot/healthbot k delete pod config-server-<pod-name>
```

The pod name is available in the output of the `/var/local/healthbot/healthbot k get pods` command.

- If you provision a segment-routed LSP by using PCEP, then the color functionality does not work. This issue occurs if the router is running on Junos OS Release 20.1R1.

Workaround: There is no known workaround.

- If you create a segment-routed LSP by using PCEP and select the routing device as **routeByDevice**, then the data displayed in **calculatedEro** is inconsistent.

Workaround: There is no known workaround.

- On the Task Scheduler page, the status of the task scheduler is not automatically updated.

Workaround: There is no known workaround.

- While adding a device collection task, by default, all devices are added in the device collection.

Workaround: There is no known workaround.

- You cannot delegate a segment-routed LSP using the GUI.

Workaround: You can delegate a segment-routed LSP through routers using the following command:

```
set protocols source-packet-routing source-routing-path {name_of_lsp} lsp-external-controller pccd
```

- While adding a network group, the Topics and Rules fields on the Add Network Group page do not display any values in the drop down. Therefore, you cannot generate any graphs on the **Monitoring > Graphs > Charts** page.

Workaround: Use **Monitoring > Graphs > Grafana** to generate graphs.

- In some systems, **check-filesystem-usage.rule** and **check-load.rule** might not display any data and therefore fail to evaluate triggers for these rules.

Workaround: There is no known workaround.

- Microservices fail to connect to PostgreSQL as PostgreSQL does not accept any connections during the primary role switchover. This is a transient state.

Workaround: Ensure that the microservices connect to PostgreSQL after the primary role switchover is complete.

- The Postgres database becomes non-operational in some systems, which leads to connection failure.

Workaround: Execute the following command in the primary node:

```
kubectl exec -n common $pod -- chmod 750 /home/postgres/pgdata/pgroot/data
```

- You can delete the default device group, **paragon-cluster**, that is configured for server monitoring.

Workaround: If you delete the **paragon-cluster** device group, then the deleted device group automatically reappears within 10 minutes. However, you need to re-instantiate any playbooks that you had instantiated prior to deleting the paragon-cluster device group.

- When you click the Help (?) icon on the top-right corner of the Paragon Automation GUI, the list of panels (Getting Started, What's New, Quick Help, About) is not displayed on the first click.

Workaround: The list of panels is displayed when you click the Help(?) icon again.

- The device discovery for Cisco IOS XR devices fails.

Workaround: Increase the SSH server rate-limit for the Cisco IOS XR device. Log in to the device in the configuration mode, and run the following command:

```
RP/0/RP0/CPU0:ios-xr(config)#ssh server rate-limit 600
```

- The default resources that are pre-installed on the Resources page might have some overlapping connections. If you add new resources there might be further overlapping connections, which reduces the effectiveness of the graphical visualization of resources.

Workaround: You have to manually realign the resources to improve the visual display of resources. After you realign, ensure that you do not refresh the Resources page; otherwise, the visualization resets to its default.

- The interactive terminal for the Pathfinder node appears blank.

Workaround: This happens when you have entered a wrong password or a username. Enter the valid username and password so that the connection is established.

- If you use BGP-LS to obtain information about the link delay and link delay variation, you cannot view the historical link delay data.

Workaround: There is no known workaround.

- The filtering option does not work on the Demands tab of the Paragon Planner Web Application GUI.

Workaround: There is no known workaround.

- The graph on the Alarms page does not reflect the latest data. That is, the graph is not updated after an alarm is no longer active.

Workaround: There is no known workaround.

- On the **Add Resources > Dependency** tab, while you are editing a new term with the dependency type as *Other Device* or *Other Network*, and if you add a second Locate Resource, the Resource Name field doesn't include a label along with the resource name.

Workaround: The label will be included along with the resource name if you perform the following steps:

1. Change the dependency type to **Local Device & Network** and revert the changes to the earlier value.
2. Update the **Label As** field in the Locate Resource section and revert the changes to the earlier value.

The Resource Name field will now include the label along with the resource name.

- In rare scenarios (For example, when Redis crashes and is auto-restarted by Kubernetes, or you have to restart the Redis server), some interfaces information is lost and interfaces are not listed on the Interface tab of the network information table. However, this issue does not affect path computation, statistics, or LSP provisioning.

Workaround: To restore interfaces in the live network model, rerun the device collection task.

- On the Tasks tab of Add New Workflow and Edit Workflow pages, the following error message is displayed when you try to edit and save an existing step without making any changes:

Name already exists

Workaround: If you have erroneously clicked the Edit option, ensure that you at least change the name of the step.

- On the Tasks tab of Add New Workflow and Edit Workflow pages:
 - Even though you click the **Cancel** option, the changes that you have made while editing a task will be saved.
 - You cannot reuse the name of a step that you have already deleted.
 - An error message will not be displayed even when you add a step with empty entries and click **Save and Deploy**.

Workaround: There is no known workaround.

- In the Paragon Automation GUI, you cannot create or delete an LSP. After you create or delete an LSP, an error message is displayed indicating a timeout for the request.

Workaround: Restart the topology server and the PCE server.

- SNMPv3 Informs are not supported in Paragon Automation Release 21.3.

Workaround: Alternatively, you can use SNMPv3 traps.

- Upgrade of some of the lower-end PTX devices with the Dual RE mode (For example, PTX5000 and PTX300) is not supported in Paragon Automation. This is because the lower-end PTX devices with the Dual RE mode do not support the bridging or bridge domain configuration.

Workaround: There is no known workaround.

- The **POST /traffic-engineering/api/topology/v2/1/rpc/diverseTreeDesign** API does not work.

Workaround: We recommend that you use the **POST /NorthStar/API/v2/tenant/1/topology/1/rpc/diverseTreeDesign** API.

- The Paragon Automation Kubernetes cluster uses self generated kubeadm-managed certificates. These certificates expire in one year after deployment unless the Kubernetes version is upgraded or the certificates are manually renewed. If the certificates expire, pods fail to come up and display bad certificate errors in the log.

Workaround: Renew the certificates manually. Perform the following steps to renew certificates:

1. Check the current certificates-expiration date by using the `kubeadm certs check-expiration` command on each primary node of your cluster.

```
root@primary1-node:~# kubeadm certs check-expiration
[check-expiration] Reading configuration from the cluster...
[check-expiration] FYI: You can look at this config file with 'kubectl -n kube-system
get cm kubeadm-config -o yaml'
```

CERTIFICATE	EXPIRES	RESIDUAL TIME	CERTIFICATE
AUTHORITY EXTERNALLY MANAGED			
admin.conf	Dec 13, 2023 13:20 UTC		
328d	no		
apiserver	Dec 13, 2023 13:20 UTC	328d	
ca	no		
apiserver-etcd-client	Dec 13, 2023 13:20 UTC	328d	etcd-
ca	no		
apiserver-kubelet-client	Dec 13, 2023 13:20 UTC	328d	
ca	no		
controller-manager.conf	Dec 13, 2023 13:20 UTC		
328d	no		
etcd-healthcheck-client	Dec 13, 2023 13:20 UTC	328d	etcd-
ca	no		
etcd-peer	Dec 13, 2023 13:20 UTC	328d	etcd-
ca	no		
etcd-server	Dec 13, 2023 13:20 UTC	328d	etcd-

ca	no			
front-proxy-client		Dec 13, 2023 13:20 UTC	328d	front-proxy-
ca	no			
scheduler.conf		Dec 13, 2023 13:20 UTC		
328d	no			
CERTIFICATE AUTHORITY	EXPIRES	RESIDUAL TIME	EXTERNALLY MANAGED	
ca	Nov 27, 2032 21:31 UTC	9y	no	
etcd-ca	Nov 27, 2032 21:31 UTC	9y	no	
front-proxy-ca	Nov 27, 2032 21:31 UTC	9y	no	

2. To renew the certificates, use the `kubeadm certs renew all` command on each primary node of your Kubernetes cluster.

```

root@primary1-node:~# kubeadm certs renew all
[renew] Reading configuration from the cluster...
[renew] FYI: You can look at this config file with 'kubectl -n kube-system get cm
kubeadm-config -o yaml'

certificate embedded in the kubeconfig file for the admin to use and for kubeadm
itself renewed
certificate for serving the Kubernetes API renewed
certificate the apiserver uses to access etcd renewed
certificate for the API server to connect to kubelet renewed
certificate embedded in the kubeconfig file for the controller manager to use
renewed
certificate for liveness probes to healthcheck etcd renewed
certificate for etcd nodes to communicate with each other renewed
certificate for serving etcd renewed
certificate for the front proxy client renewed
certificate embedded in the kubeconfig file for the scheduler manager to use
renewed

Done renewing certificates. You must restart the kube-apiserver, kube-controller-
manager, kube-scheduler and etcd, so that they can use the new certificates.

```

3. Recheck the expiration date using the `kubeadm certs check-expiration` command on each primary node of your cluster.

```

root@primary1-node:~# kubeadm certs check-expiration
[check-expiration] Reading configuration from the cluster...
[check-expiration] FYI: You can look at this config file with 'kubectl -n kube-
system get cm kubeadm-config -o yaml'
```

	CERTIFICATE	EXPIRES	RESIDUAL TIME	CERTIFICATE
AUTHORITY	EXTERNALLY MANAGED			
	admin.conf	Jan 18, 2024 21:40 UTC		
364d		no		
	apiserver	Jan 18, 2024 21:40 UTC	364d	
ca	no			
	apiserver-etcd-client	Jan 18, 2024 21:40 UTC	364d	etcd-
ca	no			
	apiserver-kubelet-client	Jan 18, 2024 21:40 UTC	364d	
ca	no			
	controller-manager.conf	Jan 18, 2024 21:40 UTC		
364d		no		
	etcd-healthcheck-client	Jan 18, 2024 21:40 UTC	364d	etcd-
ca	no			
	etcd-peer	Jan 18, 2024 21:40 UTC	364d	etcd-
ca	no			
	etcd-server	Jan 18, 2024 21:40 UTC	364d	etcd-
ca	no			
	front-proxy-client	Jan 18, 2024 21:40 UTC	364d	front-proxy-
ca	no			
	scheduler.conf	Jan 18, 2024 21:40 UTC		
364d		no		
	CERTIFICATE AUTHORITY	EXPIRES	RESIDUAL TIME	EXTERNALLY MANAGED
	ca	Nov 27, 2032 21:31 UTC	9y	no
	etcd-ca	Nov 27, 2032 21:31 UTC	9y	no
	front-proxy-ca	Nov 27, 2032 21:31 UTC	9y	no

4. Restart the following pods from any one of the primary nodes to use the new certificates.

```

root@primary1-node:~# kubectl delete pod -n kube-system -l component=kube-apiserver
root@primary1-node:~# kubectl delete pod -n kube-system -l component=kube-scheduler
root@primary1-node:~# kubectl delete pod -n kube-system -l component=kube-
```

```
controller-manager
root@primary1-node:~# kubectl delete pod -n kube-system -l component=etcd
```

Resolved Issues

This section lists the resolved issues in Juniper Paragon Automation Release 21.3.

- In Paragon Automation Release 21.3, the following Log4J vulnerability-related issues are fixed:
 - CVE-2021-44228
 - CVE-2021-4104
 - CVE-2021-45046
- The Device Group Details page does not display the entry made in the Description field while adding a Device Group.
- The status of workflows is not automatically refreshed in the Paragon Automation GUI.
- If you try to deploy any configuration during a swap of Postgres primary role, the deployment fails.
- The Bring Your Own Ingest (BYOI) feature is not supported. You cannot define your own ingest types.
- On the Device Group page, the Disable Trigger Action Schedule field is missing.
- When you run a task in a workflow, sometimes the status of the task might be displayed as Null instead of Running even though there is no error in the task. This issue does not impact the functionality, as the task status will later be updated to Completed, Failed, Error, or Pass.
- If you deploy a workflow from management daemon (mgd), then the workflow is not deployed in the GUI.
- On the **Configuration > Device** page, sometimes, the management status of a device is displayed as Down even though the connection is established.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2023 Juniper Networks, Inc. All rights reserved.