Release Notes

Published
2021-12-21

Paragon Automation, Release 21.2

JUNIPER NETWORKS | Engineering Simplicity

Table of Contents

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 like support for Juniper Networks and third-party (Cisco IOS XR) 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, 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 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 21.2. You must perform a fresh installation of Paragon Automation Release 21.2.

Licensing

In Paragon Automation Release 21.1, the licensing was honor-based. Starting from Paragon Automation Release 21.2, licenses are enforced for Paragon Insights and Paragon Pathfinder.

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

However, device licenses are not enforced in Paragon Pathfinder.

For information about Paragon Insights and Paragon Pathfinder licenses, see Paragon Insights Licensing Overview and Paragon Pathfinder Licensing Overview.

New and Changed Features

This section describes the features in each module of Juniper Paragon Automation Release 21.2.

Paragon Insights

- **Support for scheduling user-defined actions—** Paragon Automation Release 21.2 supports scheduling user-defined actions (UDAs) and notifications on the Trigger Action page. You can schedule Trigger Action schedulers to run alternatively in different instances of Paragon Automation. When your UDAs and notifications are processed in multiple parallel instances of Paragon Automation, you can minimize data loss due to node failure.

  [See Enable UDA Scheduler in Trigger Action.]

- **Support for iAgent port ingest configuration for outbound SSH connections—** Paragon Automation Release 21.2 supports iAgent port (for NETCONF connection) at the ingest level for outbound SSH traffic for all device groups. You must configure a different outbound SSH port for each device group. To avoid opening multiple ports, you can configure one port for outbound SSH traffic from all device groups.

  [See Configure Outbound SSH Port for iAgent.]

- **Support for Server Monitoring sensor—** Paragon Automation Release 21.2 supports the Server Monitoring sensor to collect data of servers and virtual machines on which the application is hosted. Server Monitoring ingest uses the third-party plug-in Node Exporter that is easy to run on servers as a binary image. The sensor collects data on CPU, memory, network, traffic, disk, and file system and writes the output to a time series database.

  [See Server Monitoring Sensor.]

- **Support for scalar fields in SNMP rules—** Paragon Automation Release 21.2 supports scalar fields along with tabular fields in SNMP rules. You can configure only scalar fields or a combination of scalar and tabular fields.

  [See Scalar Objects in SNMP.]

- **Configure and monitor action workflows (Beta)—** Paragon Automation Release 21.2 supports action-workflow monitoring (Beta). An action workflow is an action engine that you can use to configure a set of tasks. In earlier releases, you can configure user-defined actions (tasks), and cannot monitor or restart a failed action. In this release, you can configure new action workflows, monitor existing action workflows, and manage action-workflow instances on the Paragon Automation GUI.

  [See Workflow Monitoring Overview.]

- **Use Grafana UI to create graphs and visuals—**Starting with Paragon Automation Release 21.2, you can use the Grafana UI to monitor the status and health of your network devices. Grafana UI renders data from Paragon Insights time series database (TSDB). and you can view this data in the form of charts, graphs, histograms, and heat maps. The Grafana UI is an open-source visualization tool, and you can access it by clicking **Monitoring > Graphs > Grafana** in the Paragon Automation UI.

  [See Grafana Overview.]

- **Support for Paragon Insights Standard (PIN-Standard) license—**From Paragon Automation Release 21.2 onwards, Paragon Insights Standard (PIN-Standard) license replaces the HealthBot's free-tier license that was available up until the standalone HealthBot Release 3.2 and Paragon Insights 4.0.0 (Paragon Release 21.1) . With this PIN-Standard license, you can use Paragon Insights to collect telemetry data, and use default or custom playbooks without any advanced features. You can also publish telemetry data externally by using Kafka and Advanced Message Queuing Protocol (AMQP) Publish.

  [See Paragon Insights Licensing Overview.]

Paragon Pathfinder

- **Manage change requests—**Starting in Paragon Automation Release 21.2, users with the required permissions can manage change requests related to tunnels from the Change Control Management page. Use this page to authorize and track change requests to ensure that only users with the appropriate permissions can approve and deploy the changes to be implemented in the network.

  [See About the Change Management Page.]

- **Network slicing—**Starting in Paragon Automation Release 21.2, network operators can configure network slices on a network. Network slices allow services with competing requirements to be delivered over a shared infrastructure. Each slice can be configured to have a dedicated data plane to serve the slice. Additional administrative constraints can be added on the topology to steer LSP routing to that effect..

  [See Network Slicing Overview.]

- **Paragon Pathfinder uses the link statistics advertised by using BGP-LS to discover topology—**Starting in Paragon Automation Release 21.2, Paragon Pathfinder discovers the network topology and computes LSPs by using the link delay and link delay variation statistics that the routing devices advertise using BGP-LS. You can view the link delay and the link delay variation statistics as measured delay on the topology view.

  [See Configure Routers to Send Link Statistics through BGP-LS.]

Base Platform

- **User management without SMTP or e-mail verification—**Starting in Paragon Automation Release 21.2, you can add users without configuring the SMTP settings or using an activation e-mail to authenticate the new users.

  [See Add Users.]

- **Support for Junos Flex—**Paragon Automation Release 21.2 can manage devices running Junos Flex Release 20.1R3.7 or later.

  [See the table in Supported Devices and OS Versions.]

Paragon Planner

- **Store network models in Paragon Planner database—**Starting with Paragon Automation Release 21.2, you can securely store network models by saving them in an access-controlled database. Use attribute-based access control to secure the network models and retrieve them in case of an accidental data loss. In earlier releases, network models were saved on the file system without access control.

  [See Save Network Models to Paragon Planner Database.]

Graphical User Interface (GUI)

- **Support for embedded Help panels—**Starting with Paragon Automation Release 21.2, you can access the Help documentation within the user interface. Click the **Help (?)** icon to launch the following embedded Help panels:
  - Getting Started—Provides essential information to get the product up and running.
  - What's New—Provides quick access to information about new and changed features and fixed issues in the Paragon Automation release.
  - Quick Help—Provides a quick way for users to access the user guide, FAQ, and featured topics within the Paragon Automation GUI.
  - About—Provides release-specific information, including the version number and disclosures.

Paragon Installation

- **RHEL**—You can install Paragon Automation Release 21.2 on RHEL version 8 and later.
- **Docker EE**—Release 21.2 supports Docker EE version 18.03.1-ee-1 and later. If you want to use Docker EE, you must install Docker EE on all the cluster nodes before installing Paragon Automation.
- **System Requirements**—In Release 21.2, Ceph provides system storage. You need a minimum of three cluster nodes with unpartitioned disks or unformatted disk partitions for Ceph.

  [See Paragon Automation System Requirements.]

- **Open Distro**—Release 21.2 supports the Open Distro version of Elasticsearch. Open Distro provides authenticated access to the Kibana application. To log in to Kibana, you **must** configure a password in the **opendistro_es_admin_password** field in the **config.yml** file before installation.

  [See Install Paragon Automation.]

Anuta ATOM

- There are no new Anuta ATOM features in Paragon Automation Release 21.2.

Known Issues

IN THIS SECTION

This section lists the known issues in Juniper Paragon Automation Release 21.2.

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: Deleting TSDB nodes should be done 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.

General

- If you have dedicated a node for a time series database (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 unmanged 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: There is no known workaround

- The Device Group Details page does not display the entry made in the Description field while adding a Device Group.
  Workaround: Edit the device group that you created and redeploy. The entry made in the Description field is displayed.

- 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 device 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.

- After you reset the topology, if the NETCONF connection status for a node is blank, then select the node and click **More > Request NETCONF Reconnect**, otherwise, the NETCONF provisioning fails for such nodes.

  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.

- On the **Configuration > Device** page, sometimes, the management status of a device is displayed as Down even though the connection is established.

  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.

- If you try to deploy any configuration during a swap of Postgres primary role, the deployment fails.

  Workaround: Redeploy after a new Postgres primary role is elected.

- The Bring Your Own Ingest (BYOI) feature is not supported. You cannot define your own ingest types.

Workaround: The BYOI feature is available in the standalone release of Paragon Insights. To define your own ingest types, use Paragon Insights Release 4.1.0.

- Cisco 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 Device Group page, the Disable Trigger Action Schedule field is missing.

  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 the 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.

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

- If you deploy a workflow from management daemon (mgd), then the workflow is not deployed in the GUI.

  Workaround: There is no known workaround.

- 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.

- The status of workflows is not automatically refreshed in the Paragon Automation GUI.

  Workaround: Reload the page.

- 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 > Workflows > 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 PostgresSQL as PostgresSQL does not accept any connections during the primary role switchover. This is a transient state.

Workaround: Ensure that the microservices connect to PostgresSQL 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.

Resolved Issues

This section lists the resolved issues in Juniper Paragon Automation Release 21.2.

- When you use the destroy command to uninstall Paragon Automation, uninstall fails if the persistent volume that is used for backup and restore contains backup files. An error similar to the following occurs:

  ```
  TASK [local-volumes/uninstalled : Remove Bind-mounts for local-volume directories]
  ********************************************************************************************
  changed: [10.4x.xx.64] => (item=1)
  failed: [10.4x.xx.64] (item=2) => changed=false
  ansible_loop_var: item
  item: '2'
  msg: 'Error rmdir /export/local-volumes/pv*: [Errno 39] Directory not empty: ''/export/local-
  volumes/pv*'''
  changed: [10.4x.xx.64] => (item=3)
  changed: [10.4x.xx.64] => (item=4)
  changed: [10.4x.xx.64] => (item=5)
  ```

- You cannot view the statistics of segment-routing LSPs.

- While adding a maintenance event, you must not include a space in the Name field.

- If you specify an incorrect URL to access the Paragon Automation GUI, a 404 error is not displayed and you are not redirected to a known page or to an error page.

- While you reschedule the config-server microservice on another node (for example, when a node is down), if there are Kubernetes-related issues, then the Postgress database will be cleaned up and the data might not be repopulated. Due to this issue, the device groups are not listed on the Device Groups page. The Device Groups page appears to be blank.
- You cannot provision a segment-routed LSP (through PCEP or NETCONF) on Cisco IOS XR devices.
- If you change the hostname of a device on the **Configuration > Device** page or through APIs, the changes are not reflected on the Add Device Group page.
- The periodic aggregation function in the ingest pipeline ignores the packets that have arrived out of sequence (for example, this can happen for the UDP ingest). These packets are later not considered in the aggregation calculation, which in turn can result in some data deviation in the JTI telemetry data (bps/pps).
- If you deploy playbook instances back-to-back, the deployment might fail due to a database error. This is a rare scenario.
- When you update a playbook, the new changes in the playbook are not applied to the existing instances of the playbook. For example, a playbook instance that is associated with a device group is not updated when the playbook is edited or updated.
- You cannot export audit logs in PDF format.
- If you use Flex Software License Model for a device or if you are using devices running Junos OS Evolved, thenParagon Automation does not discover those devices.