

Release Notes

Published
2023-05-18

Paragon Automation, Release 22.1

Table of Contents

| | |
|---------------------------------------|----|
| Introduction to Paragon Automation | 1 |
| Installation and Upgrade Instructions | 1 |
| Licensing | 2 |
| New and Changed Features | 2 |
| Known Issues | 5 |
| Resolved Issues | 16 |

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 (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 or Release 21.2, you cannot upgrade to Paragon Automation Release 22.1. You must install Paragon Automation Release 22.1 afresh.

If your installed version is Paragon Automation Release 21.3, you can upgrade to Paragon Automation Release 22.1.

Licensing

In Paragon Insights, we've introduced the following license tiers and their related device licenses:

- 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'll receive an out-of-compliance alert in the Paragon Automation GUI if you try to deploy more devices than the number for which you've obtained licenses. However, you can continue to use the existing functionality.

You can view your license compliance status on the **Administration > License Management** page in the GUI.

In Paragon Pathfinder, we've hard-enforced the following license tiers:

- Pathfinder Standard
- Pathfinder Advanced
- Pathfinder Premium

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

Starting with Paragon Automation Release 22.1, you must upgrade the existing license key format to the new format. You can generate a new license key by using the Juniper Agile Licensing portal. For more information about generating a new license key, see [View, Add, or Delete Licenses](#).

New and Changed Features

This section describes the features in each module of Juniper Paragon Automation Release 22.1.

Paragon Insights

- **Configure pre-action and post-action tasks in rules**—Starting with Paragon Automation Release 22.1, you can configure pre-action and post-action tasks to execute action engine workflows in rules. Pre-action and post-action tasks are configurations applied to devices in device groups through playbooks. Use pre-action tasks to automate the execution of action engine workflows. With post-

action tasks, you can remove the additional configurations in devices added through the pre-action tasks.

[See [Rules Overview](#).]

- **Enhancements to ease the use of the Resources page**—Starting with Paragon Automation Release 22.1, you can perform the following additional tasks on the Resources page:
 - Upload, download, and clone resources.
 - Rearrange the sequence of the configured dependency terms. Paragon Automation executes terms in the order in which you arrange them.
 - View a visual representation of the dependency configuration before you save the configuration.
 - Use the reset icon to restore the visual panel display to the default view.
 - Change the position of resources on the visual panel and retain the changes when you navigate out of the Resources page.
 - Save filters as favorites.

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

- **Build and load BYOI custom plug-ins**—Starting with Paragon Automation Release 22.1, you can build your own ingest plugin image and load the plugin in Paragon Automation. With bring your own ingest (BYOI) custom plugins, you can reuse and analyze external telemetry data.

[See [Build and Load Custom Plugin Images](#).]

- **Track IFA devices on the Ingest Settings page**—Starting with Paragon Automation Release 22.1, you can assign the device name and device ID of the Inband Flow Analyzer (IFA) devices on the **Data Ingest > Settings > IFA Devices** page. Paragon Automation maps the device name to the respective device ID and displays the device name when you monitor the end-to-end path of IFA devices in a flow.

[See [Understand Inband Flow analyzer 2.0](#).]

- **View prepopulated graphs from Grafana UI**—Starting with Paragon Automation Release 22.1, you can view prepopulated graphs for CPU usage, disk reads, disk writes, and available memory on the Grafana dashboard. You can view a prepopulated graph of one or more nodes in a cluster.

[See [Grafana Overview](#).]

- **Display a timeline view of device and network KPI data**—Starting with Paragon Automation Release 22.1, you can display the time inspector view for **DEVICE** and **NETWORK** entity types. In earlier releases, Paragon Automation supports the time inspector view only for the **DEVICE GROUP** entity type.

[See [About the Network Health Page](#).]

Paragon Pathfinder

- **Support for segment routing with IPv6 as the forwarding plane**—Starting with Paragon Automation Release 22.1, you can provision segment routing tunnels that use IPv6 as the forwarding plane (also called SRv6 tunnels).

NOTE: Currently, Paragon Pathfinder does not support binding segment identifiers (SIDs) and anycast SIDs for SRv6 tunnels. In addition, closed-loop telemetry-driven use cases are not supported.

[See [Segment Routing Overview](#).]

- **Support for configuring path computation timeout**—Starting with Paragon Automation Release 22.1, during stateful path computation for PCEP request (PCReq), PCEP reply (PCReply), and REST API, the path computation state is held for the specified timeout (rounded to the next multiple of 5).

If the user request or the PCC provisions the policy within the timeout, the already computed state and bandwidth are used. If the timeout expires before the user or the PCC provisions the policy, a new path is computed. After the PCC sends the PCReport (PCRpt), the LSP is reported with the computed Explicit Route Object (ERO).

Default: 10 seconds.

You can configure a *path-computation-state-timeout* from the **Pathfinder Settings** page (**Configuration > Pathfinder Settings > Path Computation Server**).

[See [Modify Pathfinder Settings from the GUI](#).]

Base Platform

There are no new features related to the base platform in Paragon Automation Release 22.1.

Paragon Planner

There are no new features related to Paragon Planner in Paragon Automation Release 22.1.

NOTE: Paragon Planner Web Application is a beta feature in Paragon Automation Release 22.1.

Paragon Installation and Upgrade

- **Support for air-gap installation on nodes with Red Hat Enterprise Linux base OS**—Starting with Paragon Automation Release 22.1, you can use the air-gap method to install Paragon Automation on nodes with the Red Hat Enterprise Linux base OS. In the air-gap method, the cluster nodes don't need Internet access during installation. Paragon Automation Release 22.1 is qualified to work with RHEL version 8.4.

[See [Air-Gap Install Paragon Automation on RHEL](#).]

Anuta ATOM

There are no new features related to Anuta ATOM in Paragon Automation Release 22.1.

Known Issues

IN THIS SECTION

- [Installation | 5](#)
- [General | 7](#)

This section lists the known issues in Juniper Paragon Automation Release 22.1

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.

- 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.

- In a multi-primary node setup, if the primary nodes use DHCP for their IP address configuration and if the DHCP must be renewed, then Calico may use an incorrect IP address for the IP-IP tunnel.

Workaround: Configure the primary nodes using a static IP address.

- While installing Paragon Automation using the air-gap method, the following error occurs:

```
task path: /runtime/roles/docker/tasks/prepare-redhat.yml:40
fatal: [ppflabapp102]: FAILED! =>
  msg: |-
    The task includes an option with an undefined variable. The error was: list object has no
    element 0
```

Workaround: Edit the following configuration variables in the `config-dir/config.yml` file and then install Paragon Automation using the air-gap method:

```
docker_version: '20.10.13-3'
containerd_version_redhat: '1.5.10-3'
```

- The Paragon Automation Setup installation bundle lists Foghorn as one of the components. The installation fails if you select Foghorn.

Workaround: Foghorn is not supported in Release 22.1. While installing Paragon Automation, we recommend that you do not select the Foghorn option.

General

- From Junos OS Release 22.4R1 onwards, the multipath feature for PCEP segment routing-traffic engineering (SR-TE) is enabled, by default. However, if the PCEP peer is not multipath capable, PCEP sessions are not established.

Workaround: To ensure that PCEP sessions are established, disable the PCEP multipath feature using the following command.

```
set protocols pcep disable-multipath-capability
```

- 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 the topology service fails to connect to influxdb, it retries connecting to influxdb continuously leading to high CPU utilization. The topology service fails to connect to influxdb, as it tries to connect to localhost by default, instead of `tsdb.healthbot`.

Workaround: Delete the topology service and the topology-api service using the following procedure.

Prerequisites

- Ensure that the `curl` command is installed on any one of the primary nodes.
- Ensure that you have the required privileges to execute Kubernetes-specific commands such as `kubectl -n healthbot get pods`. Set the `KUBECONFIG` variable corresponding to your installation, if required.
- Ensure that all the pods in the `healthbot` namespace are running and their status is healthy.

Perform the following steps to delete the topology service and the topology-api service.

1. Log in to one of the primary nodes of your Kubernetes cluster.

2. Change directory to **healthbot**.

```
cd /var/local/healthbot
```

3. Use a text editor to remove the following lines from the **healthbot.sys** file under the 'DEFAULT-SERVICES' section.

```
TOPOLOGY = topology
TOPOLOGY_API = topology-api
```

4. Run the `curl` command to sync the **healthbot.sys** file across all the nodes.

```
curl --location --request POST http://localhost:7005/publish \
  --header 'Content-Type: application/json' \
  --data-raw '{
    "channel": "common",
    "files": [
      "/var/local/healthbot/healthbot.sys"
    ],
    "recursive": true,
    "notify": true,
    "prune": false,
    "delete": false
  }'
```

5. Run the following command to remove the topology service and the topology-api service.

```
kubectl -n healthbot delete deployments topology topology-api
```

- 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`.

- While adding a device, if you specify a source IP address that is already used in a network, you may not be able to add the device to a device group, deploy a playbook, encounter function ingest-related errors, and so on.

Workaround: Fix the conflicting source IP address. Click the Deployment Status icon and commit the changes.

- 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.

- 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.

- 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: For the procedure to enable the topology filter, see the [Enable the Topology Filter Service](#) topic.

- 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 have enabled the outbound SSH at a device group-level, you cannot disable the outbound SSH for one of the devices in the device group.

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.

- 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: Once you provision the path, the path respects the slice ID constraints and the path appears correctly in the path preview.

- 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: Upgrade the Junos OS to Release 21.4R1.

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

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:

```
for pod in atom-db-{0..2}; do
  kubectl exec -n common $pod -- chmod 750 /home/postgres/pgdata/pgroot/data
done
```

- 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
```

- 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.

- 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:
 - 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.

- 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.

- Paragon Automation doesn't show alarms for Nokia devices.

Workaround: There is no known workaround.

- While configuring an SRv6 LSP with the routing method as routeByDevice, you must specify a value for the segment routing-Explicit Route object (SR-ERO); otherwise, you cannot use the SRv6 LSP to carry traffic.

Workaround: While adding a tunnel, on the Path tab, add hops to specify the required or preferred routing type.

- If a device-controlled SRv6 LSP is discovered from the network, the path highlighted for this LSP will be incorrect irrespective of whether or not you specify an Explicit Route object (ERO) for the route.

Workaround: There is no known workaround.

- Sometimes, you may not be able to delete segment routing LSPs in bulk.

Workaround: You can force delete the LSPs that are not deleted during the process of bulk deletion.

- 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.

- In the Paragon Automation GUI, on the Tasks tab of the 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.

- The PCEP session is displayed as Down if you restart the PCE server.

Workaround: Restart the topology server by using the `kubectl delete pods ns-toposerver-<POD_ID> -n northstar` command.

- You cannot provision a secondary LSP for an existing LSP. This is applicable for both PCE-initiated and NETCONF LSPs.

Workaround: Use the routeByDevice option as the routing method for the secondary LSP.

- On the Administration > License Management page, you cannot view the SKU name of a license when you select the license and then select More > Details.

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.

- When you configure the outbound SSH for iAgent, the data for the configured rule will not be generated.

Workaround: There is no known workaround.

- A zero percent value of packet loss is displayed between the links if you have configured Two-Way Active Management Protocol (TWAMP). This is incorrect because TWAMP does not support exporting packet loss for IS-IS traffic engineering.

Workaround: There is no known workaround.

- If you are using a device with MPC10+ line cards and if the device is running on a Junos OS Release other than Release 21.3R2-S2 or Release 21.4R2-S1, then the statistics for logical interfaces are not collected. However, the statistics for physical interfaces and LSPs are collected.

Workaround: Upgrade the Junos OS release to Release 21.3R2-S2 or 21.4R2-S1. Also, ensure that you have upgraded Paragon Automation to Release 22.1.

- While adding a P2MP group, you cannot choose a specific Z-endpoint as the **Node Z List** field lists only the **Select All** option.

Workaround: Select the **Select All** option and later remove the unwanted Z-endpoints.

- When you undelegate an LSP, the LSP status is displayed as delegated. When you try to undelegate the LSP again, the router configuration might be modified to add explicit route objects (ERO).

Workaround: Refresh the Tunnel tab before you undelegate the LSP again.

- Paragon Pathfinder does not bring down a delegated SR LSP when the SR LSP does not meet slice constraints if the SR LSP's status is locally routed.
- If you create a topology group with slice ID greater or equal to $2^{**}32$, the topology group ID will not match the slice ID.
- 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 | | | |
| 364d | admin.conf | Jan 18, 2024 21:40 UTC | | |
| | | no | | |
| ca | apiserver | Jan 18, 2024 21:40 UTC | 364d | |
| | | no | | |
| ca | apiserver-etcd-client | Jan 18, 2024 21:40 UTC | 364d | etcd- |
| | | no | | |
| ca | apiserver-kubelet-client | Jan 18, 2024 21:40 UTC | 364d | |
| | | no | | |
| 364d | controller-manager.conf | Jan 18, 2024 21:40 UTC | | |
| | | no | | |
| ca | etcd-healthcheck-client | Jan 18, 2024 21:40 UTC | 364d | etcd- |
| | | no | | |
| ca | etcd-peer | Jan 18, 2024 21:40 UTC | 364d | etcd- |
| | | no | | |
| ca | etcd-server | Jan 18, 2024 21:40 UTC | 364d | etcd- |
| | | no | | |
| ca | front-proxy-client | Jan 18, 2024 21:40 UTC | 364d | front-proxy- |
| | | no | | |
| 364d | scheduler.conf | Jan 18, 2024 21:40 UTC | | |
| | | 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 22.1.

- On the Add Filter page, the Save option is not displayed if you create more than one filter on the GUI pages.
- 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.
- The interactive terminal for the Pathfinder node appears blank.
- 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.
- SNMPv3 Informs are not supported in Paragon Automation Release 21.3.
- 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.
- 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.

- 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.
- 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.