

# Release Notes

Published  
2023-10-26

## Paragon Automation, Release 23.1

---

# 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	17

# 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). In addition to providing base platform capabilities, Paragon Automation offers 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.

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 22.1, you can upgrade to Paragon Automation Release 23.1.

If your installed version is Paragon Automation Release 21.3, we recommend that you upgrade to Release 22.1 and then upgrade that to Release 23.1.

If your installed version is Paragon Automation Release 21.2, you cannot upgrade to Paragon Automation Release 23.1. You must install Paragon Automation Release 23.1 afresh.

# 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](#).

If you have a license key that was generated for a version of Paragon Automation earlier than Release 22.1 you will need to upgrade the license key format to the new format before you can install it in Paragon Automaton Release 23.1. 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 23.1.

### Paragon Pathfinder

- **Audit logging enhancements**—Starting with Paragon Automation Release 23.1, you can view the Paragon Pathfinder user operations on the **Audit Logs** page. User operations include options to add, modify, or delete network elements such as nodes and links. See [Audit Logs Overview](#).

- **Import topology filter rules**—Starting with Paragon Automation Release 23.1, you can import multiple topology filter rules and apply the rules to your network. Use the **Topology Filter** page to import the required rules as comma-separated values (CSV) file that contains the rules. See [About the Topology Filter Page](#).
- **Disaster recovery support**—Starting with Paragon Automation Release 23.1, you can deploy Paragon Automation at two different geographical locations to ease disaster recovery. When the Paragon Pathfinder component is down in one location, the Paragon Pathfinder component at the other location continues to manage the Path Computation Client (PCC)-delegated label-switched paths (LSPs) in your network.

You must perform the following tasks on the two deployments:

- Configure network discovery by using BGP Link State (BGP-LS) and Path Computation Element Protocol (PCEP)
- Configure analytics
- Independently onboard the device, schedule playbooks, and perform other deployment tasks.

The federated exchange that you configure on the messaging bus synchronizes the two Paragon Automation deployments to:

- Manage topology-related changes [for example, add, modify, and delete nodes, links, and point-to-point label-switched paths (P2P LSPs)]
- Optimize LSPs
- Reset and synchronize topology
- Add or remove LSP delegation

For information on disaster recovery, see [Disaster Recovery Overview](#) and for information on configuring disaster recovery, see [Configure Disaster Recovery for Paragon Pathfinder](#).

- **Support for collecting colored segment routing LSP statistics on Cisco Devices**— Starting with Paragon Automation Release 23.1, you can collect colored Segment Routing LSP (Label Switched Path) statistics for Cisco devices. A new rule controller-telemetry /ctrl-label-switched-path-sr-color-snmp-cisco-aggregation is added as an optional rule. As Cisco devices do not send color attributes over PCEP, you have to manually patch the color attribute for colored Segment Routing LSPs or run the device collection on the device to collect the LSP statistics. For more information, see [Collect Analytics Data Overview](#).

## Base Platform

- **Authenticate a device using RADIUS credentials**—Starting with Paragon Automation Release 23.1, you can manage your network devices by using RADIUS authentication. You can configure RADIUS credentials to discover a device and manage all the connections to the device.

In earlier releases, after you discover a device by using the device credentials, Paragon Automation uses the automatically generated credentials for all the connections to the device. In Release 23.1, you can discover and manage your network devices by using RADIUS credentials as well as credentials generated by Paragon Automation. For more information, see [Add Devices](#) and [Edit Devices](#).

## Paragon Insights

We've not added any new features related to Paragon Insights in Paragon Automation Release 23.1.

## Paragon Planner

We've not added any new features related to Paragon Planner in Paragon Automation Release 23.1.

**NOTE:** Paragon Planner Web Application is a beta feature in Paragon Automation Release 23.1.

## Paragon Installation and Upgrade

- **Configure external registry**—Starting with Paragon Automation Release 23.1, you can configure the Docker registry on an external node. You can also use custom external registries in place of the Paragon Automation-generated registries. In earlier releases, Paragon Automation generates the Docker registry and stores it in Ceph storage. In Release 23.1, you can configure the following registry options:

- An existing external custom registry
- One Paragon Automation-generated registry on an external node
- Multiple Paragon Automation-generated registries on multiple external nodes

All registry nodes must be in the same subnet so that you can connect to them by using a virtual IP (VIP) address.

We've updated the **inventory** and **config.yml** files to include the external registry configuration parameters in Release 23.1. For more information on configuring an external registry, see [Configure External Docker Registry](#).

- **Schedule configuration backup**—Starting with Paragon Automation Release 23.1, you can schedule a back up of your Paragon Automation configuration using a backup script. Restoring your backed up configuration is a manual process. However, you can use the restore script to view a sequential list of commands used to restore your backed up configuration. For more information, see [Backup and Restore Scripts](#).

## Anuta ATOM

Integration of Paragon Automation Release 23.1 with Anuta ATOM has not been tested. We recommend that you validate the integration before upgrading to newer releases of Anuta ATOM.

# Known Issues

## IN THIS SECTION

- [Installation | 5](#)
- [General | 6](#)

This section lists the known issues in Juniper Paragon Automation Release 23.1

## Installation

- When you provision virtual machines (VMs) on VMware ESXi servers, if you add the block storage disk before adding the disk with the base OS, Ceph sometimes incorrectly identifies the drives and creates the cluster using the incorrect drive, resulting in the base OS being destroyed.

Workaround: Add the first disk as the base OS (larger drive) and then add the smaller block storage disk.

- 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.

- 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'
```

## General

- The root cause analysis (RCA) feature is disabled, by default, in Paragon Automation Release 23.1.

Workaround: To enable the RCA feature, perform the following steps.

1. Ensure that you set the `kubeconfig` environment variable or the `~/.kube/admin.conf` file is present.
2. Ensure that the Paragon Insights services are up and running. To verify that the status of the pods is Running, use one of the following commands.

```
root@primary-node:~# kubectl -n healthbot get pods
```

or

```
root@primary-node:~# kubectl get po -n healthbot
```

3. Log in to one of the primary nodes.
4. Change the working directory to `/var/local/healthbot`.

```
root@primary-node:~# cd /var/local/healthbot
```

5. Use a text editor to edit the `healthbot.sys` file and change the value for `ENABLE_RCA` to "True" in the inference-engine section.

```
...
ENABLE_RCA = True
```

Save and close the file.

6. Publish the changes to all the other nodes in the cluster using the following command.

```
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
}'
```

7. Restart the inference-engine, alerta, and config-server services using the following commands.

```
root@primary-node:~# /var/local/healthbot/healthbot restart -s inference-engine --device-group healthbot
```

```
root@primary-node:~# /var/local/healthbot/healthbot restart -s alerta --device-group healthbot

root@primary-node:~# /var/local/healthbot/healthbot restart -s config-server --device-group healthbot
```

## 8. Verify that all the services are up and running.

```
root@primary-node:~# kubectl -n healthbot get pods
```

**NOTE:** This procedure enables the feature at the global level. You can continue to use the **Root Cause Analysis** option on the **Configuration > Device Groups** and **Configuration > Network** pages, to enable or disable the RCA feature at the device group level.

- Junos OS Release 22.4R1 and later have a limitation with SR-TE LSPs.

For PCEP sessions to be established, you must disable the multipath feature using the following command:

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

Secondary path is not supported.

- Admin group of an SR-TE LSP learned from PCEP disappears after topology synchronization, if the LSP has configured state.

Workaround: Modify SR-TE LSP to persist the admin group learned from PCEP.

- Safe mode status is always false when ns-web pod starts.

Workaround: None.

- “Delegation Bit” column in the tunnel table is always empty.

Workaround: View delegation bit by selecting the LSP, right-click and select **Show Details**.

- Old messages in queue are being processed after federation link is recovered.

Workaround: None.

- You get an incorrect safe mode status after you the modify source-of-truth flag during safe mode.

Workaround: None.

- Demand is missing from the Network Archive, in the Java Client, but the demand traffic is available.

Workaround: None.

- Sometimes devices with NETCONF disabled appear with NETCONF status Up.

Workaround: Edit the device profile without any changes to trigger reloading of device profile.

- Color for SR-TE LSPs originating from Cisco IOS-XR devices is visible only if the LSP is initially discovered from device collection.

Workaround: None.

- Config viewer in the Integrity Check tab does not work. You can use the config viewer available in Topology tab instead.

Workaround: None.

- Toposerver saves the live properties (ero and rro) path for each LSP twice.

Workaround: None.

- The P2MP logical tree-view can incorrectly show nodes multiple times resulting in cycles in a tree. This is due to the issue where toposerver saves live properties path for each LSP twice.

Workaround: None.

- The maintenance mode simulation feature and the corresponding maintenance reports are not available.

Workaround: Run failure simulation in Java Planner.

- Paragon Automation uses only the Network Configuration Protocol (NETCONF) and not the Path Computation Element Protocol (PCEP) to collect point-to-multipoint (P2MP) status for Cisco routers. You must select only **NETCONF** as the provisioning method when you create P2MP groups for Cisco routers.

Workaround: None.

- You cannot disable the source-of-truth flag when the deployment is in safe mode.

Workaround: Restart the toposerver pod to disable source-of-truth flag during safe mode.

- When you select multiple delegated label-switched paths (LSPs) belonging to a single ingress router and click **Return Delegation to PCC**, only one of the LSPs becomes device controlled. An issue in Junos causes this scenario.

Workaround: Select one LSP at a time and click **Return Delegation to PCC** individually for each LSP.

- The operational status of a delegated SR-TE LSP remains down after its destination node is rediscovered.

Workaround: You must sync the network model after the delegated SR-TE LSP destination node is rediscovered.

- PCE server is unable to reconnect to rabbitmq after rabbitmq is restarted.

Workaround: Restart the ns-pceserver pod.

- You cannot modify the use-federated-exchange setting from REST API/UI.

Workaround: Modify the use-federated-exchange setting directly from the cMGD CLI and restart the toposerver for the change to take effect.

- 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 on the device 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`.

- 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: None.

- 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 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: None.

- 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.

- 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.

- If you have enabled outbound SSH at a device group-level, you cannot disable 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.

- 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: None.

- 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: None.

- 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: None.

- 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: None.

- 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: None.

- 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.

- 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 sometimes displayed as Down if you restart all pods in the northstar namespace.

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

- 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: None.

- 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: None.

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

Workaround: None.

- 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: None.

- 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 23.1.

- 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 23.1.

- 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.
- 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.
- 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.
- P2MP groups configured by Path Computation Element Protocol (PCEP) with flowspec mapping to multicast VPN service is not supported.  
  
Workaround: There is no known workaround.
- On the Task Scheduler page, the status of the task scheduler is not automatically updated.

Workaround: None.

---

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.