

Release Notes

Published
2024-07-24

Paragon Automation, Release 24.1

Software Highlights

- Support for RHEL 8.10
- Ability for non-root users to run commands in the paragon CLI utility
- Ability to provision segment routing policies on Cisco IOS XR devices using NETCONF

Table of Contents

[Introduction | 1](#)

[Installation and Upgrade Instructions | 1](#)

[Licensing | 2](#)

[New and Changed Features | 2](#)

[Deprecated Features | 3](#)

[Known Issues | 4](#)

[Resolved Issues | 18](#)

Introduction

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 new 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](#).

Use these release notes to find new and updated features, software limitations, and open issues in Paragon Automation Release 24.1.

Installation and Upgrade Instructions

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

NOTE:

You can directly upgrade only from Paragon Automation Release 23.2 to Release 24.1. If your release is earlier than Release 23.2, you must install Release 24.1 afresh. However, in order to migrate your current release configuration to Release 24.1, you can use the back up and restore functionality. For more information on upgrade, see [Upgrade to Paragon Automation Release 24.1](#).

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](#).

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 24.1. 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 24.1.

Paragon Installation and Upgrade

- **Red Hat Enterprise Linux (RHEL) 8.10**—Paragon Automation Release 24.1 is qualified to work with RHEL 8.10.

[See [Installation Prerequisites on Red Hat Enterprise Linux](#).]

- **Run paragon CLI utility commands as a non-root user**—Starting in Paragon Automation Release 24.1, a non-root user with superuser (sudo) privileges can run the paragon CLI utility commands to analyze, query, and debug the Paragon Automation setup.

[See [Troubleshoot Using the paragon CLI Utility](#).]

Paragon Pathfinder

- **Provision segment routing policies on Cisco IOS XR devices**—Starting in Paragon Automation Release 24.1, you can provision segment routing policies on Cisco IOS XR devices by using NETCONF as the provisioning method.

Base Platform

We've not added any new features related to the base platform in Paragon Automation Release 24.1.

Paragon Insights

We've not added any new features related to Paragon Insights in Paragon Automation Release 24.1.

Paragon Planner

We've not added any new features related to Paragon Planner in Paragon Automation Release 24.1.

NOTE: Paragon Planner Web Application is a beta feature in Paragon Automation Release 24.1.

Deprecated Features

This section lists the features that are deprecated or for which support is withdrawn from Paragon Automaton Release 24.1.

- **Grafana UI**

You cannot access the Grafana UI from Paragon Automation. To access the Grafana UI, you must:

1. Install Grafana.

See [Grafana Documentation](#) for more information.

2. Expose the TSDB port by running the `/var/local/healthbot/healthbot tsdb start-services` command.

NOTE: In Paragon Automation, the TSDB port is not exposed by default. To use external tools such as Grafana, you need to run a query to the TSDB directly (and not through APIs) to expose the TSDB port.

For more information, see [Backup and Restore the TSDB](#).

- **Charts**

Known Issues

IN THIS SECTION

- [Installation | 4](#)
- [General | 6](#)

This section lists the known issues in Juniper Paragon Automation Release 24.1

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 `deploy-federated-exchange` command output displays that installation has failed when you configure disaster recovery in a dual cluster deployment. You can ignore the failure message but you *must* execute the following command on all the primary nodes of both the clusters:

```
kubectl patch deployment ns-web -n northstar --type "json" -p '[{"op": "add", "path": "/spec/
template/spec/containers/0/env/-", "value": {"name": "USE_FEDERATED_EXCHANGE", "value":
"true"}}]'
```

Workaround: None.

- When a failure impacts both diverse pair of LSPs, the Path Computation Server (PCS) will not route the LSPs along the lesser diversity level path or along the non-diverse path. The LSPs are not routed until the PCS can find a path matching the configured diversity level.

Workaround: None

- When a failure impacts both diverse pair of LSPs, the Path Computation Server (PCS) will not route the LSPs along the non-diverse path. The LSPs are not routed until the PCS can find a path matching the configured diversity level.

Workaround: Remove and reapply the diversity group.

- The minimum variation threshold value under bandwidth sizing settings of a container subLSP is shown as 0 despite configuring it in the container. Under normal conditions, there is no impact to bandwidth sizing of the subLSP as the bandwidth sizing task fetches this value from the container instead of the subLSP. However, in certain scenarios, it is possible that the subLSP could be resized to new bandwidth value when the configured minimum variation threshold has not been breached.

For more details on this issue, contact Juniper Networks Technical Assistance Center (JTAC).

- During bandwidth sizing, an active secondary LSP that has bandwidth sizing enabled might not get resized. When this issue occurs, the RSVP utilization of the links in secondary path could be incorrectly updated.

Workaround: None.

- Making changes to Paragon Pathfinder settings (**Configuration > Network Settings**) by using the UI might require more than one attempt for the modification to take effect. You might have to click **Save** more than once.

Workaround: The same changes could be made by using cMGD CLI which is accessible from the master node running the `pf-cmgd` command.

- Under certain conditions during container normalization, one or more container subLSPs that were supposed to be removed will continue to remain. These container subLSPs will remain in the network as independent LSPs not associated with the container. A mismatch in the number of subLSPs of a container stated in the subLSPs column under Container LSP tab and actual number of LSPs having the name of the container as prefix under Tunnel tab, could be considered as an indication of this problem.

For more details on this issue, contact Juniper Networks Technical Assistance Center (JTAC).

- Container LSP can be configured with bandwidth sizing settings that are inherited by its subLSPs. Under certain circumstances, when a user disables bandwidth sizing option in the container after having enabled it in the past, it does not get disabled in the existing subLSPs.

Workaround: None.

- Manual reprovisioning of a subLSP of container would lead to addition of data to the LSP object. As a result, the following issues could occur:
 - If the container is bandwidth sizing-enabled and a non-zero minimum variation threshold is configured, the specific subLSP could get resized despite the traffic through the subLSP not exceeding its signaled bandwidth by at least minimum variation threshold value.
 - The subLSP could end have different bandwidth sizing settings than the container if the container bandwidth sizing settings are modified later.
 - Failure in removal of subLSP during container normalization when the bandwidth falls below merging bandwidth.

For more details on this issue and on instructions to remove the additional data that gets added to the internal state, contact Juniper Networks Technical Assistance Center (JTAC).

- Under certain scenarios such as container normalization failure on lack of available paths, additional internal state would get added to container subLSP objects that could lead to the following issues:
 - If the container is bandwidth sizing-enabled and a non-zero minimum variation threshold is configured, the specific subLSP could get resized despite the traffic through the subLSP not exceeding its signaled bandwidth by at least minimum variation threshold value.
 - The subLSP could end have different bandwidth sizing settings than the container if the container bandwidth sizing settings are modified later.
 - Failure in removal of subLSP during container normalization when the bandwidth falls below merging bandwidth.

For more details on this issue and on instructions to remove the additional data that gets added to the internal state, contact Juniper Networks Technical Assistance Center (JTAC).

- When one or more nodes in a working Kubernetes cluster is unavailable, it could result in the following unexpected behavior:
 - PCEP status of all nodes shown as down though the PCEP connection status is up on the router.
 - Network topology not displayed in the UI.

For more details on this issue, contact Juniper Networks Technical Assistance Center (JTAC).

- Paragon Pathfinder could compute a path that violates the maximum hop constraint that is configured in the tunnel. These scenarios explain how the maximum hop constraint is violated:
 - When the Path Computation Server (PCS) restarts, the down LSP is provisioned without considering the maximum hop constraint.
 - During network failure, the LSP is rerouted without considering maximum hop constraint.
 - During path optimization, the LSP is optimized without considering maximum hop constraint.

Workaround: Use the reprovision option if an alternate path that does not violate the configured constraint is available.

- The path computed by Paragon Pathfinder for a standby LSP with maximum hop constraint could violate the configured constraint.

Workaround: None.

- There is a possibility that the PCS is unable to find LSP with link diversity in a topology that has multiple parallel links between nodes.

Workaround: None.

- When PCEP session is disabled, the LSP operational status will move to *unknown state* after running device collection.

Workaround: None.

- A link might go missing when creating a network archive task.

Workaround: Create a new network archive task.

- Unable to route VPN demand due to issues in the network.

Workaround: None.

- Alarms do not respond when Cisco devices are initially configured with NETCONF on port 22.

Workaround: Modify the NETCONF port on your Cisco device to *<new-port-no>* and ensure that the changes are saved. After this, revert the port settings back to port 22.

- When you add multicast demands in the GUI, the node Z field is empty.

Workaround: None.

- When you add multiple new tunnels, the traffic values from previously deleted tunnels (that were cached) are displayed.

Workaround: None.

- When you add new diverse tunnels, sometimes the traffic values from previously deleted tunnels (that were cached) are displayed.

Workaround: None.

- The Toposerver does not clear or update the topology after losing connection to the BMP pod.

Workaround: None.

- When a link is down, Paragon Pathfinder does not reroute a delegated SR LSP with preferred Explicit Route Object (ERO) and **Route By Device** routing method.

Workaround: Use the **Default** routing method.

- If you run a simulation directly after you perform a Diverse Multicast Tree Design, the report in **Tunnel Traffic on Links (Tunnel Layer Simulation Report > Peak Network Statistics)** is incorrect.

Workaround: Save the network after you perform a Diverse Multicast Tree Design and close it. Re-open the network and then run the simulation.

- While simulating failure scenarios (**Tools > Options > Failure Simulation**), if you run a multiple failure simulation first and then run a single failure simulation after, the report in **Tunnel Traffic on Links (Tunnel Layer Simulation Report > Peak Network Statistics)** is incorrect. The report displays multiple failure simulation values instead of single failure.

Workaround: De-select all options on the **Multiple Failure** tab before simulating a single failure scenario.

- Link Utilization simulation report might show negative values during the double failure scenario.

Workaround: None.

- When a device hostname is changed, the change is not reflected across all databases.

Workaround: Perform the following steps so that the new device hostname is reflected across all databases and components.

1. Before the hostname change, remove the device from all the device groups (controller or other playbooks).

2. Ensure that the device references are deleted from all the different Paragon Automation components. Navigate to the **Configuration > Devices** page.

- a. Select the device.
- b. Click the trash can icon to delete the device. The **Delete Device** page appears.
- c. Select **Force Delete** and click **Yes**.

3. Onboard the device again, using the device onboarding workflow from the **Configuration > Devices** page.

The device should now be onboarded with new hostname. The device properties especially system-id (important for receiving the JTI streams) should also be updated.

4. Add the device with the new hostname back to the device groups.
5. (Optional) Verify all the device stats in Influxdb using Grafana or on the device CLI. The database should be updated with the new hostname.

- The Network Configuration Protocol (NETCONF) provisioning method for point-to-multipoint (P2MP) LSPs is not supported in Cisco IOS-XR routers.
- On Cisco IOS-XR routers, the P2MP sub-LSP status is not supported in configuration state for CLI-provisioned P2MP LSPs.

Workaround: None.

- 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.

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

Workaround: Set the federation link queue expiry time close to Toposerver federation link failure detection time (default is 3*5s).

- You cannot use the NETCONF and Path Computation Element Protocol (PCEP) methods to provision P2MP LSPs for Cisco IOS-XR routers using the Paragon Automation UI.

Workaround. Provision the P2MP LSPs using the CLI. After the configuration is parsed, run a Device Collection task to view the LSPs.

- 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`.

- When you undelegate a delegated LSP, the planned bandwidth of the LSP is based on the bandwidth reported by the device instead of the user input value.

Workaround: None.

- 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.

- 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 Grafana. To log in to Grafana or ESDB, you must configure a password in the **grafana_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
kubect1 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
admin.conf          Jan 18, 2024 21:40 UTC
364d                no

```

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

- 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 24.1

- Symmetric pair LSPs might not be routed symmetrically upon threshold crossing rerouting.

Workaround: None.

- Traffic charts are now supported for devices with dual Routing Engines which are onboarded with re0 or re1 suffixed to their hostnames. However, graphs are only supported if the hostname-suffixes are in lower case and in the -re0 or -re1 format. For example: vmx101-re0 or vmx101-re1

Workaround: None

- Controller Sites are not included in network archive for Paragon Planner.

Workaround: None.

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

Workaround: None.

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

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.

- 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.

- LSPs on the optimal path may receive unnecessary PCEP update during PCS optimization.

Workaround: None.

- An error in the diagnostics (**Configuration > Data Ingest > Diagnostics > Application**) feature causes application tests to fail.

Workaround: None.

- On the **Network > Topology > Tunnel** tab, when you hover over the Filter (funnel) icon and select **Add Filter**, the **Add Criteria** page is displayed. If you select **Color** in the **Field** list, the field value is displayed as **plannedProperties** instead of Color.

Workaround: None.

- Path analysis report is empty.

Workaround: Run a device collection task before performing path analysis. Note that, Path analysis report can be empty if LSPs are already on optimal path.