Cisco Crosswork Network Controller 7.2 Closed-Loop Network Automation

First Published: 2026-01-31

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

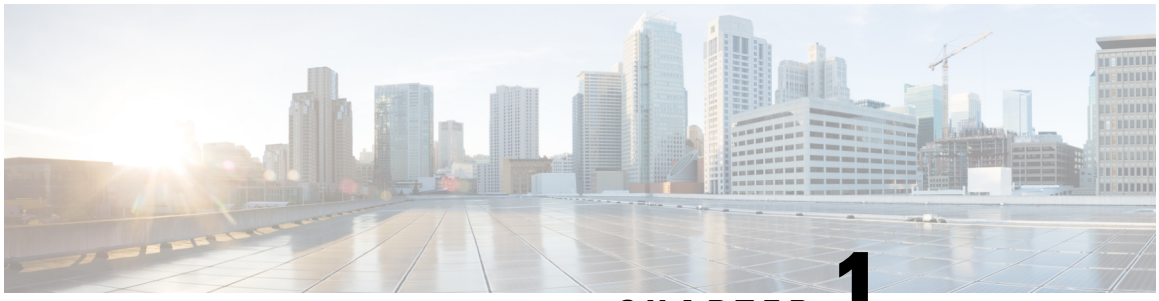

CHAPTER 1

Overview

This section contains the following topics:

Audience

This guide describes the steps required to begin using Change Automation and Health Insights after installation. It is intended for experienced network administrators who want to use these components in their network. Before proceeding, ensure that you are familiar with the following topics:

- Networking technologies and protocols, such as IS-IS, BGP, and other relevant protocols.
- Network monitoring and troubleshooting
- Platform Infrastructure and Crosswork Network Controller components: For additional details on installation and setup, see the Cisco Crosswork Network Controller Installation Guide.

Overview of Change Automation and Health Insights

Health Insights and Change Automation are optional components that can be installed with the Crosswork Network Controller (available in Crosswork Essentials, Crosswork Advantage, or Crosswork Premier). These components are included in the Add-on package available on Cisco.com.

Note The package is available only for existing users of Crosswork Network Controller Add-on components.

The components provide a ready-to-use solution supporting the following use cases:

- Monitor Key Performance Indicators (KPIs) and notify of any anomalies.
- Prepare network changes triggered by changes in KPIs and roll out these changes.
- Automate change and remediation.

Change Automation

Change Automation helps to codify workflows using parameterized plays and stitches them into playbooks for execution. It offers a collection of plays and playbooks designed to help you easily implement changes to the network for various situations. You also have the option to create your own playbooks to simplify the network operations or to implement network configurations in response to changing circumstances identified through Health Insights KPIs.

Health Insights

Health Insights offers real-time, telemetry-based Key Performance Indicator (KPI) monitoring and intelligent alerting. The alerts are based on predefined templates or user-defined logic. These alerts can be tied to the playbooks to implement closed-loop automation workflows.

Health Insights supports building KPIs based on telemetry using MDT, SNMP, or GNMI. The collected data is evaluated in one of the following four possible ways (using UI based tools):

- No alert
- Standard deviation
- Two-level threshold
- Rate change

Other configurations are also possible using the Cisco Crosswork APIs. For more details, see Cisco Crosswork Network Automation APIs.

Cisco Crosswork API

All the Crosswork Network Controller components provide a robust set of APIs that allow it to be integrated with other tools you use to manage and configure your network. For more details on the product APIs, see the Cisco Crosswork Network Controller API Documentation on Cisco DevNet.

Integration with other Cisco and non-Cisco products

Change Automation and Health Insights support a wide range of use cases. These capabilities can be further extended by integrating with other Cisco and non-Cisco products. Sample configurations for integration with various tools are available on Cisco DevNet. For customers requiring additional customization, Cisco CX Services offers specialized engagements to tailor the solution to specific operational needs.

The following products can be integrated to extend the functionality of Change Automation and Health Insights:

- **Cisco Crosswork Planning**: Crosswork Planning provides traffic and topology analysis to Change Automation and Health Insights. It gives a cross-sectional view of traffic, topology, and equipment state. For more information, see Cisco Crosswork Planning.
- **Cisco Network Services Orchestrator (NSO)**: Cisco Network Services Orchestrator acts as the default provider to configure the devices according to their expected functions, including configuring any required model-driven telemetry (MDT) sensor paths for data collection. Cisco Network Services Orchestrator is vital in supplying device management and configuration-maintenance services. For more information, see Network Services Orchestrator (NSO).

- **Optimization Engine**: Optimization Engine provides real-time network optimization. Some plays enable integration with Crosswork Optimization Engine so that the optimization decision is based on the KPIs being tracked in Health Insights. For more information, see Cisco Crosswork Optimization Engine Data Sheet.
- **Cisco Element Management Functions (EMF)**: A library of functions that provides detailed device inventory, software image management, device alarm, device key metrics, configuration related functions and ZTP.
- **Non-Cisco products**: Change Automation and Health Insights supports the loading of models for non-Cisco equipment which will enable the creation of KPIs and in some cases, the execution of plays. For more information on how to do these advanced integrations, see the Crosswork Network Controller Administration Guide and the Crosswork Network Controller API Documentation on Cisco DevNet. If you require assistance with these integration efforts, contact your account team.

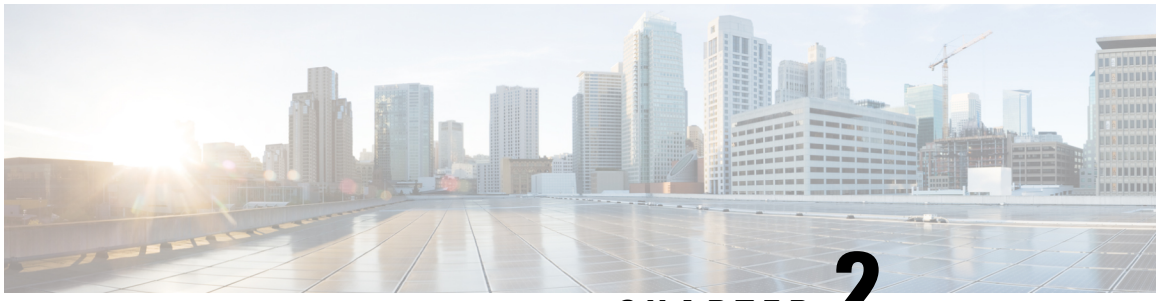

CHAPTER 2

Get Started

This section contains the key workflows and an overview of Change Automation and Health Insights dashboard:

Getting started

After installing the Change Automation and Health Insights applications, there are a few initial steps that administrators must complete to ensure the application is properly configured and ready for use. For installation details, see the Cisco Crosswork Network Controller Installation guide.

Change Automation can be used independently or as part of workflows that leverage Health Insights or other applications.

Before you begin:

- Make sure to install the Change Automation and Health Insights applications. See the Cisco Crosswork Network Controller Installation guide.

This initial setup includes configuring system settings, assigning appropriate user access levels, creating device access groups, and device tags. The following sections provide detailed guidance on each of these steps.

Verify installation and configure system settings

This section describes the system settings that must be configured to begin using Change Automation.

Change Automation provides several ways to run playbooks. As part of the initial activation process, you will need to select your preferred method for running playbooks.

- Manually ("on demand") or via scheduled execution. These two methods are typically used for playbooks that accomplish data collection, configuration changes, or SMU deployment independent of any KPI-related fault detected in the network.
- Manually or automatically when the playbook is tied to a KPI. These methods are typically used when you want to run a playbook intended to remediate a fault detected in the network. Key parameters needed to run the playbook are populated when the alert tied to the KPI is triggered.

To verify application installation and configure system settings:

1. Navigate to **Network automation** > **Dashboard**.

   The first time you access Change Automation, you will see the **Welcome to network automation** page where the system will prompt you to complete the application's configuration.

   Figure 1: Welcome to network automation

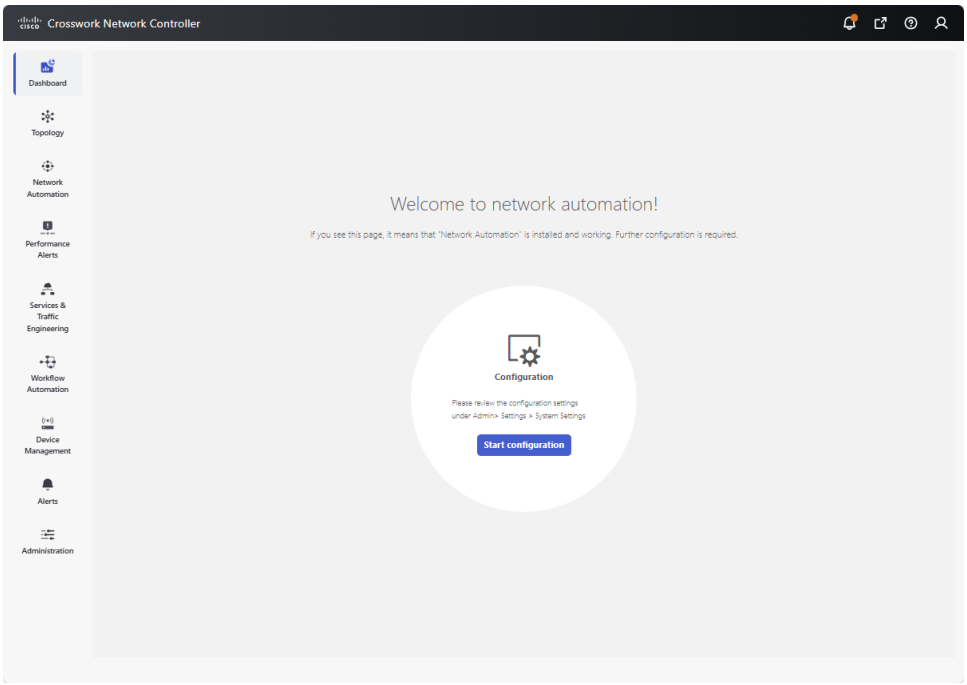

2. Click **Start configuration** or optionally, navigate to **Administration** > **Settings** > **System settings** > **Network automation** > **Device override credentials** to make the settings.

Figure 2: System settings - Device override credentials

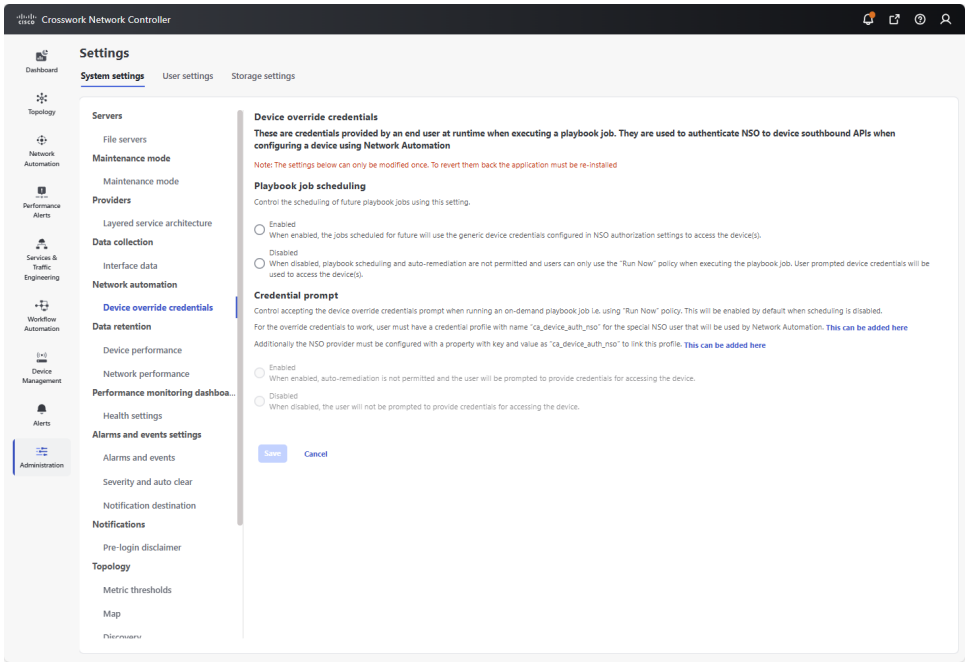

3. During this initial setup, you will be required to configure device override credential settings for Playbook job scheduling and Credential prompt. These settings work together to determine whether playbooks can be run unattended, either as a scheduled task or as automatic remediation to changes detected in the network by Health Insights.

Note The Change Automation settings can only be configured once. If you need to modify them later, Change Automation must be re-installed. Before re-installing, export any plays or playbooks you have created, and after re-installing, import them. For more information, see Export plays, on page 26, Import custom plays, on page 26, Export playbooks, on page 34, and Import playbooks, on page 35.

Playbook job scheduling: This setting enables or disables the ability to schedule playbook jobs.

Credential prompt: This setting determines whether user interaction is required when running an on-demand playbook.

Note In addition to the administrative user account needed to allow Crosswork Network Controller to communicate with NSO, a second credential profile with name "ca_device_auth_nso" must be created for the override credentials to work when running playbooks. For details on creating credential profile in NSO, see Cisco Crosswork Change Automation NSO Function Pack Installation Guide.

If Playbook job scheduling is...	and the Credential prompt is...	This will...
enabled	**disabled**	enable both automatic playbook scheduling and auto-remediation. With this setting: • future jobs will use generic device credentials configured in Cisco NSO to access the devices. • you will not be prompted to provide credentials for accessing the device. • the system will be able to automatically detect network issues and take corrective actions without requiring manual intervention. • you can run playbooks automatically whenever the KPI linked to that playbook raises an alert of sufficient severity.
enabled	**enabled**	only enable automatic playbook scheduling. With this setting: • automatic remediation is not permitted. • you will be prompted to provide credentials for accessing the device. • future jobs will use generic device credentials configured in Cisco NSO to access the devices. • you can run playbooks automatically whenever the KPI linked to that playbook raises an alert of sufficient severity.
disabled	**enabled** (by default)	disable both automatic playbook scheduling and auto-remediation. With this setting: • you can only use **Run Now** to run the playbook job. • you will be prompted to provide credentials for accessing the device. • credential prompt is enabled by default and cannot be disabled.

Special considerations

- If **Credential prompt** is **enabled**: While executing Device Config plays, entering incorrect device override credentials will cause the playbook execution to fail. However, for a Check play or Data Collection play, the device override credentials are not validated and the playbook will execute successfully irrespective of their accuracy.
- If **Credential prompt** is **disabled**: Only user IDs with write permissions for **Administration APIs** under **Change Automation** can complete the credential profile and provider setup tasks. If you are unsure if your user ID has the required privileges, you can check by selecting **Administration** > **Users and Roles** > **Roles** and inspecting the ID's privileges.

4. Once you have made your selections as per your requirements, click **Save** to commit to these settings.

Assign user access levels

Once the system settings are configured, an administrator should review user roles to ensure that all users have the appropriate level of access required for Change Automation and Health Insights management.

Change Automation API permissions allow users to run, import, and create plays and playbooks. Only users with write permissions for **Administration APIs** can disable or enable playbook execution access and assign labels.

Health Insights KPI Management APIs allow users to:

- Create, edit, or delete KPIs and KPI profiles.
- Monitor the status of KPI-related jobs.
- Configure alerts to proactively manage network performance.

To enable user access for Change Automation and Health Insights:

Procedure

Step 1 Go to **Administration** > **Users and Roles** > **Roles**.

Step 2 Under the **Roles** pane, select the role to which you want to grant access.

Step 3 Under **Global API permissions**, find **Change Automation**.

Step 4 For **Play** and **Playbook** APIs, enable **Read** and **Write** check boxes (as necessary).

Figure 3: Global API permissions - Change Automation

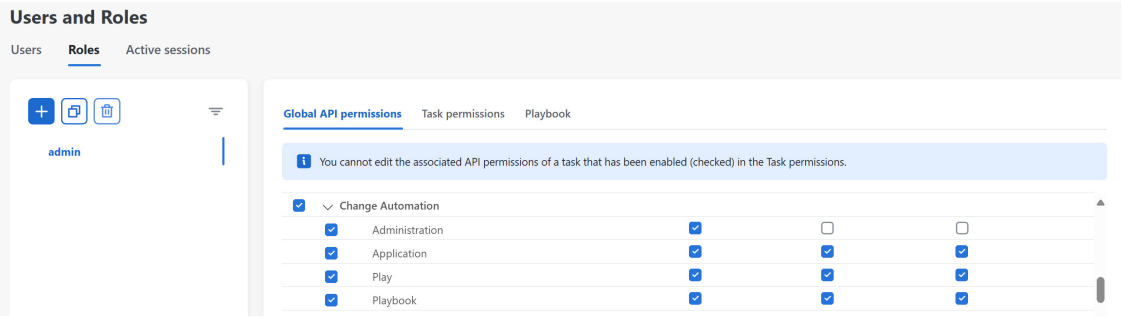

Step 5 Scroll down to find **Health Insights**.

Step 6 For **Health Insight KPI Management**, enable **Read** and **Write** check boxes (as necessary).

Figure 4: Global API permissions - Health Insights

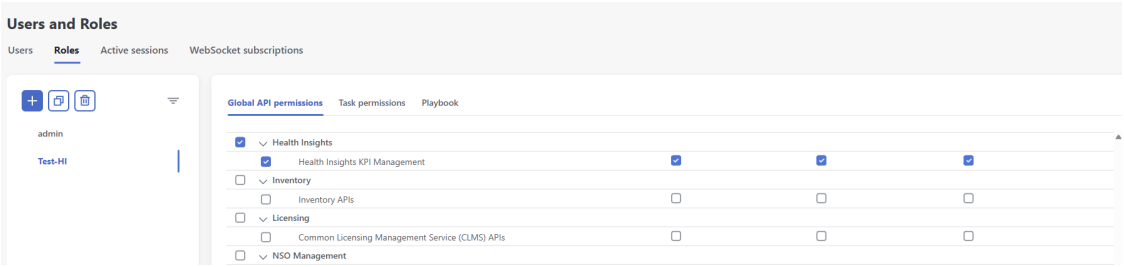

Step 7 Click **Save**.

Create Device Access Groups

After administrators have assigned the required permissions to users, they can further control which devices users are allowed to access and manage. This can be achieved using **Device Access Groups**, which logically group devices to streamline access control. Non-administrative users with the system-level task of Device Access Groups management can create and manage these groups as needed.

For more information on Device Access Groups and how to associate a user with a Device Access Group, see the **Manage Device Access Groups** section in the Cisco Crosswork Network Controller Administration guide.

Create device tags

Device tags in Crosswork Network Controller are optional but highly beneficial for organizing and streamlining network management. They allow you to group devices based on shared attributes, making it easier to run playbooks or create KPI profiles. Tags can also provide useful information, such as a device's physical location or its administrator's email ID. Once you have a set of devices with the same tag, you can use it to:

- run playbooks on a specific group of devices. For example, if devices are tagged by region, you can run a playbook exclusively on devices in that region.
- build KPI profiles tailored to a group of devices. Newly tagged devices are automatically included in the existing KPI profiles, eliminating the need for manual updates.

For more information on creating device tags, see the **Manage Tags** section in the Cisco Crosswork Network Controller Administration guide.

Using Change Automation

This section outlines various workflow scenarios that illustrate how to effectively use Change Automation. These scenarios demonstrate only a fraction of the capabilities of the Crosswork Network Controller and highlight the platform flexibility. By using these concepts and examples, you can build a virtually unlimited combination of tools to meet your unique operational needs.

Run playbooks manually

The workflow below describes the steps to run playbooks manually.

Workflow	Action
1. Select the playbook you want to run. You can choose to run it manually from a list of devices or by leveraging device tags. You also specify the appropriate run-time parameters that you want the playbook to apply.	See About running playbooks, on page 36
2. Select the playbook execution mode, which determines whether the playbook is being tested and validated or actively executed to gather data or make changes to the device.	See: • Playbook execution order, on page 37 • Perform a dry run of a playbook, on page 38 • Run playbooks in single stepping mode, on page 44 • Run playbooks in continuous mode, on page 50
3. You can link any Health Insights KPI to your playbook. The playbook will automatically run whenever the linked KPI triggers an alert in response to events like threshold crossings, topology changes, or flapping conditions.	See: Link a playbook to a Health Insights triggered KPIs, on page 12

Schedule playbooks

The workflow below describes the steps to automate routine network tasks and verify that each routine change is completed correctly.

Note This workflow is applicable only if scheduling is enabled in the Change Automation settings. For more information, see Verify installation and configure system settings, on page 5.

Workflow	Action
1. Identify routine maintenance tasks (such as throughput checks, software upgrades, SMU installs, and so on) that you perform on a regular schedule, and that may be suitable for automation using one or more Change Automation playbooks.	See About running playbooks, on page 36 and View the playbook list, on page 19.
2. Configure playbooks to perform these tasks at the desired time. You can choose to run it manually from a list of devices or by leveraging device tags. You also specify the appropriate parameters that you want the playbook to apply and select the execution mode.	See About running playbooks, on page 36 and Schedule playbook runs, on page 56.

Workflow	Action
3. Review the Change Automation job history to review the current status of the playbook. If the job fails, the details will be available.	See Use the change automation dashboard, on page 59 and View or abort playbook jobs, on page 58.
4. If you prefer to run the playbook manually on a scheduled basis, that's perfectly fine. However, if you have identified scenarios where you would like it to run automatically, you can link the playbook to the KPIs. The playbook will automatically run whenever the linked KPI triggers an alert.	See: Link a playbook to a Health Insights triggered KPIs, on page 12

Develop custom playbooks

The following workflow will enable you to develop a Change Automation custom play or playbook.

Workflow	Action
1. Review the existing plays and playbooks to see if they fully or partially meet your needs.	From the main menu, choose **Network Automation** > **Play List** or **Playbook List**.
2. If required, build new plays and then a new playbook with new or existing plays, as necessary, to meet your requirements.	See About custom plays, on page 21 and About customizing playbooks, on page 28.
3. Once you have build your playbook, you can decide to run it manually or schedule it to run as an automated routine task.	See: • Run playbooks manually, on page 11 • Schedule playbooks, on page 11
4. If you prefer to run the playbook manually on a scheduled basis, that's perfectly fine. However, if you have identified scenarios where you would like it to run automatically, you can link the playbook to the KPIs. The playbook will automatically run whenever the linked KPI triggers an alert.	See: Link a playbook to a Health Insights triggered KPIs, on page 12

Link a playbook to a Health Insights triggered KPIs

The following workflow describes the steps to link playbooks to KPIs and run them automatically or as needed.

Workflow	Action
Run playbook manually: Manual playbook execution allows you to maintain human oversight and ensures that the planned remediation effort effectively addresses the issue. It is particularly useful in situations where remediation may require prior notifications or scheduling during a maintenance window. For frequently triggered KPIs with a known remediation playbook, link the playbook to the KPI and run it manually.	See: Link KPIs to playbooks and run them manually, on page 75 Use the Remediation icon shown in View alerts for network devices to trigger a run of a linked playbook from a device or KPI alert.
Run playbook automatically: Automatically running playbooks for KPIs with known remediation eliminates the need for human intervention, ensuring a prompt resolution. For frequently triggered KPIs with a known remediation playbook and minimal risk of unintended consequences, link the playbook to the KPI and configure it to run automatically. See Verify installation and configure system settings, on page 5 to ensure your setup is properly configured to enable automatic playbook execution.	See Link KPIs to playbooks and run them automatically, on page 77

Using Health Insights

Before you begin:

1. Confirm if the Yang modules provided include the data point you want to evaluate. If the Yang module contains the data you need, then review whether one of the four available KPI templates are adequate to collect data and evaluate the data point.
   - If yes, build a new KPI profile using the existing template.
   - If no, proceed to the next step
2. If the Yang module contains the data point but there isn't an existing KPI template to evaluate it, then build a new KPI with the tools available in the developer network (developer.cisco.com).
3. If the Yang module does **not** include the data point you need:
   - Get the new Yang module that includes the required data point.
   - Load it on the data collection UI.
   - Build a new KPI.

Monitor device KPIs using Health Insights

The following table describes the steps to monitor device KPIs using Health Insights application.

Workflow	Actions
1. Plan which Cisco-supplied KPIs you want to begin using based on each device's function and the device performance characteristics you want to monitor.	See List of Health Insights KPIs, on page 64. To create a new KPI that fits your requirements, see Create a new KPI, on page 73.
2. Based on your experience or by using the recommendation engine, group the KPIs to form KPI profiles.	See Create a new KPI profile, on page 80.
3. Enable the appropriate KPI profiles on the devices you want to monitor.	See Enable KPI profiles on devices, on page 83.
4. Make sure that the collections are provisioned on the device (MDT collections).	See Verify the deployment status of enabled KPIs, on page 86

Develop custom KPIs

The following workflow describes the steps to determine whether developing custom Health Insights KPIs is necessary for your specific requirements and provides guidance on how to proceed if you choose to create them.

Workflow	Action
1. Review the existing KPIs to ensure the telemetry you want to monitor is not already available.	See List of Health Insights KPIs, on page 64.
2. Review the data available from the devices you want to monitor to see if they can supply the needed information: • If they can, proceed with building a custom KPI. • If they cannot, we must load a new Yang module.	See Create a new KPI, on page 73.
3. Determine if the Yang module we have provided includes the data point you wish to evaluate. If it does, determine whether one of the available KPI templates can evaluate it. If it can, proceed with building a new KPI. If not, you must build the KPI with the tools available in the Cisco DevNet and then import it into Crosswork Network Controller. Once you import the KPI, you can add it to your profile.	
4. Build the custom KPI and add it to a KPI profile.	See Create a new KPI, on page 73 and Create a new KPI profile, on page 80.
5. Enable the new KPI profile on a test device.	See Enable KPI profiles on devices, on page 83.
6. Confirm that collections are working.	

Workflow	Action
7. Confirm that the data reported matches your expectations and, if necessary, investigate the alarms raised by the new KPI. Be aware that KPIs that depend on data over time to establish baseline performance will need some time to establish a baseline before they provide meaningful data.	See View alerts for network devices, on page 90.
8. If the KPI profile meets expectations, enable it on all devices where applicable. **Warning** When enabling KPI profiles on many devices, ensure that sufficient capacity is available on Data Gateway. If adequate capacity is not available and if you enable the KPI profiles on a large number of devices, it may cause overload and outage. To check Data Gateway load, see *Health Insights CDG load calculator* at Cisco Crosswork Network Automation APIs.	See Enable KPI profiles on devices, on page 83.
9. Make sure the KPI profile was deployed on the device (MDT only) and that the collection jobs are functioning.	See Verify the deployment status of enabled KPIs, on page 86.

Closed-loop automation

The following workflow describes the steps to follow when using Health Insights to run a remediation playbook from Change Automation in response to the performance challenges detected in the network by a KPI. A remediation playbook can be:

- Linked to a KPI, alerting the operator to run the playbook and make the remediation easier.
- Linked to a KPI and selected for automatic execution without operator intervention.

Step	Action
1. Research the KPIs that are triggering alerts and determine the best corrective action to take for the situation your network has experienced.	Follow the instructions in Monitor Network Health and KPIs, on page 63, using the View alerts for network devices, on page 90 to research the alerts and their possible causes.
2. Review the plays and playbooks to determine which will best address the alerting KPI. For example: • Look for an existing playbook that could resolve the issue. • Look for existing plays that could be combined to resolve the issue. Create a new playbook with those plays.	See Create a custom play using templates, on page 21 and Create a custom playbook through the UI, on page 28.

Step	Action
3. Try out the selected playbooks and see if they are applicable to your purposes. As you experiment, adjust the playbook parameters as needed.	See: Perform a dry run of a playbook, on page 38 Run playbooks in single stepping mode, on page 44 Run playbooks in continuous mode, on page 50
4. If required, build new plays and then build new playbooks with the combination of plays needed to make the desired change(s) to the network.	See Create a custom play using templates, on page 21 and Create a custom playbook through the UI, on page 28.
5. (Optional) Run playbook manually - Manual playbook execution allows you to maintain human oversight and ensures that the planned remediation effort effectively addresses the issue. It is particularly useful in situations where remediation may require prior notifications or scheduling during a maintenance window. For frequently triggered KPIs with a known remediation playbook, link the playbook to the KPI and run it manually.	See Link KPIs to playbooks and run them manually, on page 75. Use the Remediation icon shown in View alerts for network devices, on page 90 to trigger a run of a linked playbook from a device or KPI alert.
6. (Optional) Run playbook automatically - Automatically running playbooks for KPIs with known remediation eliminates the need for human intervention, ensuring a prompt resolution. For frequently triggered KPIs with a known remediation playbook and minimal risk of unintended consequences, link the playbook to the KPI and configure it to run automatically. See Verify installation and configure system settings, on page 5 to ensure your setup is properly configured to enable automatic playbook execution.	See Link KPIs to playbooks and run them automatically, on page 77

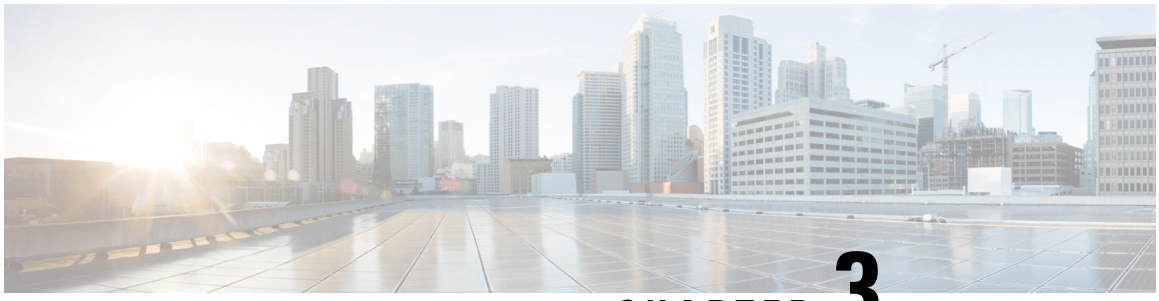

CHAPTER 3

Automate Network Changes

This section contains the following topics:

Change Automation overview

The Change Automation application automates the process of deploying changes to the network. You can define automation tasks to achieve the intended network states in Change Automation using playbooks that consist of plays written using YAML. You can then push configuration changes to Cisco Network Service Orchestrator (NSO), which deploys these changes to the network devices.

Change Automation, in conjunction with health insights, allows operators to build automation in a *closed-loop framework*. Changes are deployed to the router or other device using programmable APIs, and the intent of the change is verified using telemetry that comes back from the router. Change Automation relies on telemetry to verify the intent of the change, avoiding the need to frequently poll the device for updates.

The following is a high-level Change Automation workflow:

1. Review the existing plays and playbooks to see if they fully or partially meet your needs.

Note Change Automation comes with a robust library of playbooks, each with its own collection of configuration and check plays.

2. Build playbook as required:
   - If the required playbook is available, use it.
   - If some combination of existing plays accomplishes the task, build a new playbook using those plays.
   - If some of the required plays are not available, create new plays and build a new playbook using the new and existing plays.

3. Dry run the playbook to test if it performs as expected.
4. Deploy the playbook.

Change Automation allows you to customize and generate plays and playbooks using its API interface. For more information, see About custom plays, on page 21 and About customizing playbooks, on page 28.

Configure Change Automation settings

Configuring system settings is a post-installation activity and is the first task to be performed after installing Change Automation.

For more information, see Verify installation and configure system settings.

View the play list

The **Play List** window of the Change Automation application gives you a consolidated list of all the plays in the system.

From the main menu, select **Network Automation** > **Play List** to view the **Play List** window.

Figure 5: Play List

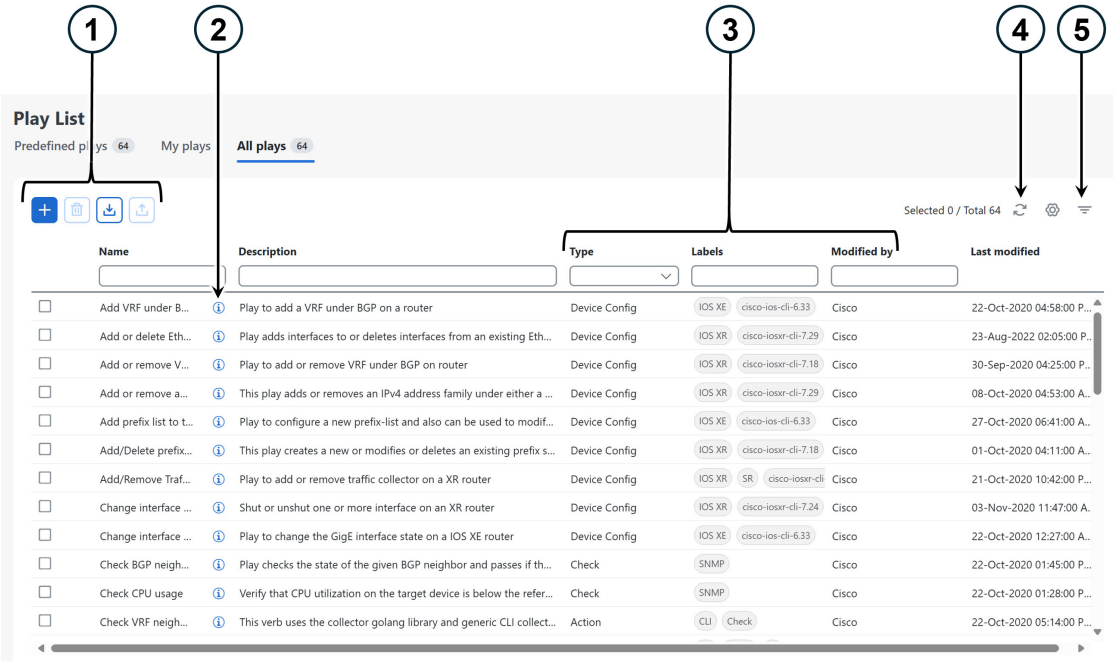

Item	Description
1	Click to create a custom play. See Create a custom play using templates, on page 21.
	Click to delete a custom play. See Delete custom plays, on page 27.
	Click to import a custom play from a gzipped TAR archive file. See Import custom plays, on page 26.
	Click to export a custom play as a gzipped TAR archive file. See Export plays, on page 26.
2	Click to see a pop-up **Play details** window showing the play's description and schema. When you are finished viewing these details, click × or the **Close** button to close the pop-up window.
3	The **Type** column indicates the type of the play. You can click the column headings (Name, Description, Type, Labels, and Modified by) to sort the table using that column's data.
4	Click to refresh the plays list.
5	Click to set filter criteria on one or more columns in the table.
	Click to clear any filter criteria you may have set.

View the playbook list

The Change Automation application's **Playbook List** window (in the following figure) gives you a consolidated list of all the playbooks in the system. To view the **Playbook List** window, select **Network Automation** > **Playbook List**.

Figure 6: Playbook List

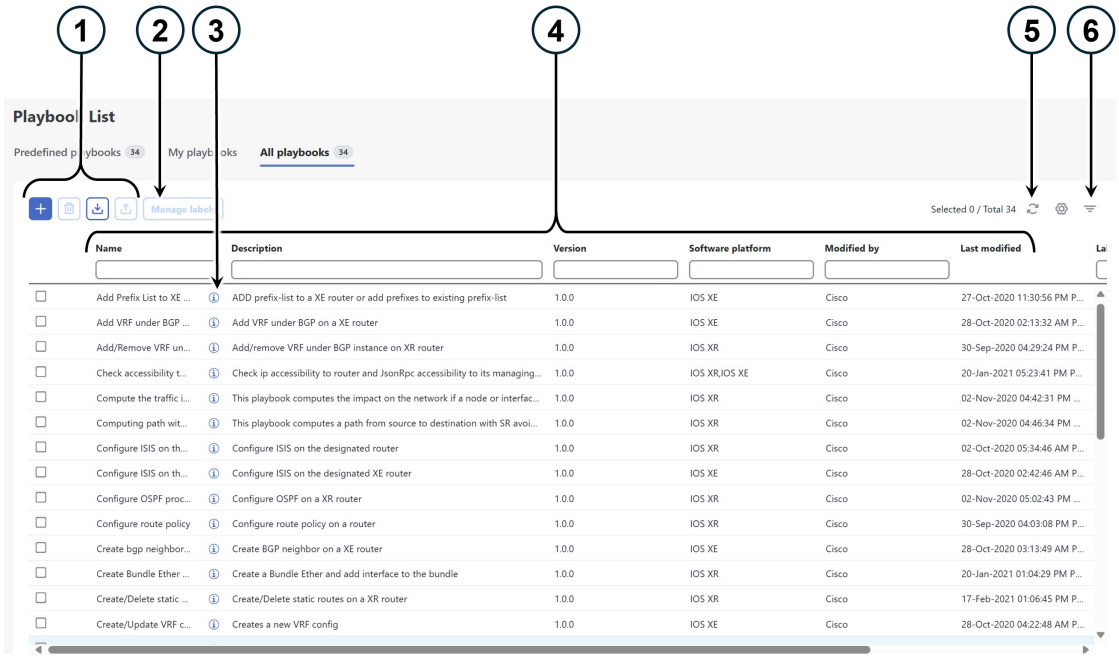

Item	Description
1	Click + to create a custom playbook. See Create a custom playbook through the UI, on page 28.
	Click 🗑 to delete the currently selected custom playbook. See Delete custom playbooks, on page 35.
	Click ⬇ to import playbooks from a gzipped TAR archive file. See Import playbooks, on page 35.
	Click ⬆ to export the currently selected playbook(s) as a gzipped TAR archive file. See Export playbooks, on page 34.
2	Click **Manage labels** to assign a label(s) to the playbook. Assigning label(s) to the playbooks allows the system administrator to control which playbooks each user role is allowed to run.
3	Click ⓘ to see a pop-up **Playbook details** window showing the playbook's description, software compatibility, version number, and its plays. When you are finished viewing these details, click ✕ or the **Close** button to close the pop-up window.
4	Click the **Name**, **Description**, **Version**, **Software platform**, and **Last modified** column headings in the table to sort the table by that column's data. You can also choose which columns are shown and set quick or advanced filters for any column.
5	Click ↻ to refresh the playbooks list.

Item	Description
6	Click to set filter criteria on one or more columns in the table.
	Click to clear any filter criteria you may have set.

About custom plays

Change Automation allows you to create your own custom plays, either based on Cisco models or from scratch. You can also import, export, and delete your custom plays.

You can create custom plays in any of the following types:

- **Check play**: Verifies the data from your devices using a logical expression.
- **Data collection play**: Collects data from your devices.
- **Device config play**: Performs configuration changes on your device
- **Service play**: Provisions and manages a service that is deployed.

Note You cannot edit, export, or delete Cisco-supplied plays.

Note Check play and Data collection play supports MDT and SNMP collection.

Create a custom play using templates

This section explains the procedure to create a custom play. The stages of play creation vary depending on the play type you choose:

- **Check play**: *Select play type > Select sensor path > Build check expression > Review play*
- **Data collection play**: *Select play type > Select sensor path > Build filter expression > Review play*
- **Device config play** or **Service play**: *Select play type > Configure play (using sample payload in JSON format) > Review play*

Procedure

Step 1 From the main menu, choose **Network Automation** > **Play List**. The **Play List** window is displayed.

Step 2 Click to create a custom play. The **Select play type** window opens displaying the types of plays supported and a description for each. The stages of creation are also displayed, and it varies depending on the play type you select.

Figure 7: Select play type

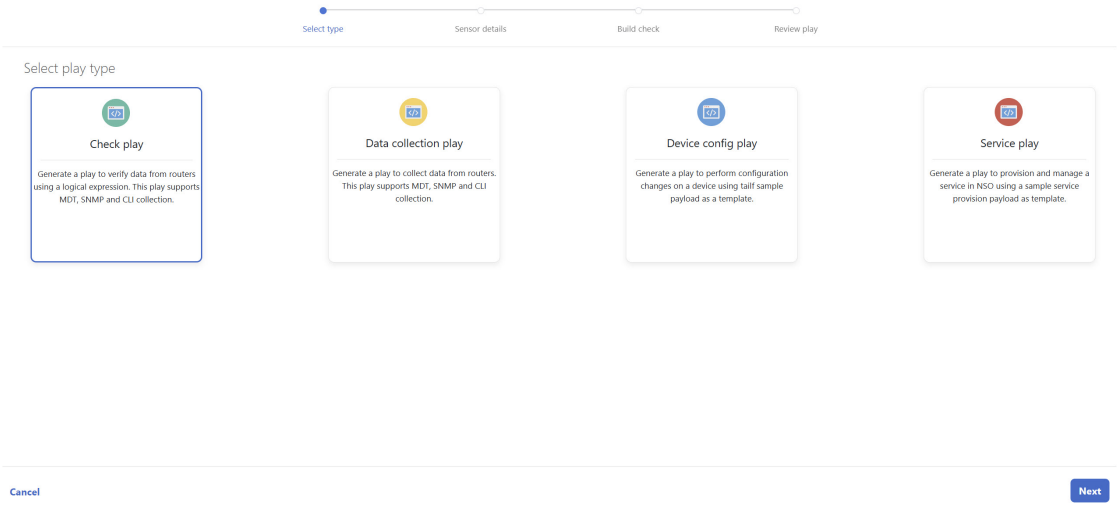

Select the play type that you want to create and click **Next**.

Step 3 Creating a Check play or Data collection play

When creating Check or Data collection plays, Cisco provides YANG modules for Cisco products. The process that is described in this section assumes that the sensor that you want to use or the field that you want to modify is included in the modules that are provided by Cisco. If the sensor or field is not listed in the default YANG modules, Cisco allows you to expand the device coverage. For information on loading a new or modified module, see Device package management in the *Cisco Crosswork Network Controller Administration* guide.

a) In the **Select sensor paths** window, select the required YANG module, Gather path, and Sensor paths. Click **Next** to continue.

Figure 8: Select sensor paths

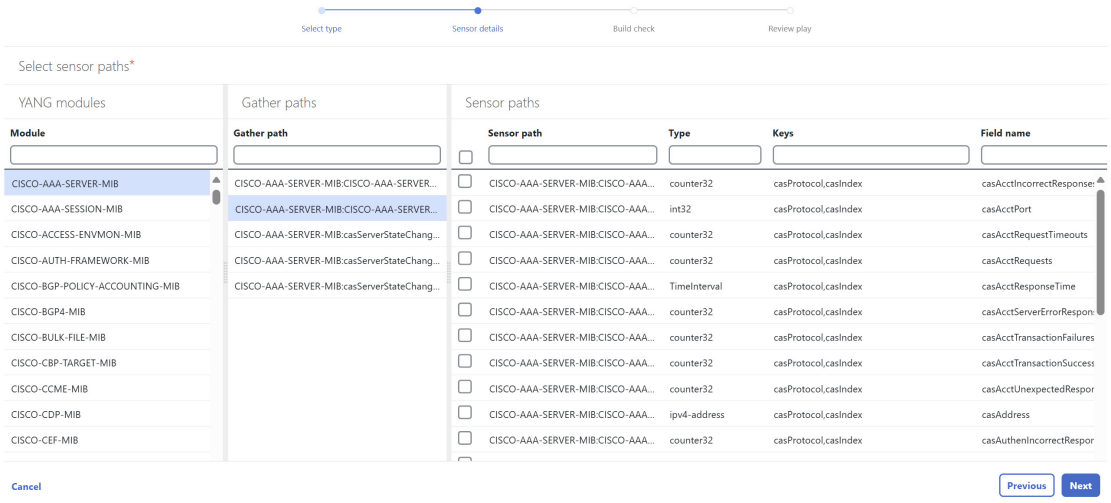

b) Depending on the play type you have selected, you must **Build check** (for Check play) or **Build filter** (for Data Collection play) to apply in your play. Click **Add rule** to add a logic expression using the keys and fields of the

selected sensor path(s). Click **Add group** to add a new logic group. Select the sensor field, operator, and value from the drop-down lists. Select the desired logic operation (AND/OR) between each rule or group.

Click the **Runtime** check box if you prefer to enter the value of the sensor field dynamically during run time. If you select this check box, the *value* field is disabled, and you will be prompted to enter the input parameter when this play is executed (as part of a playbook) during run time.

Figure 9: Check expression

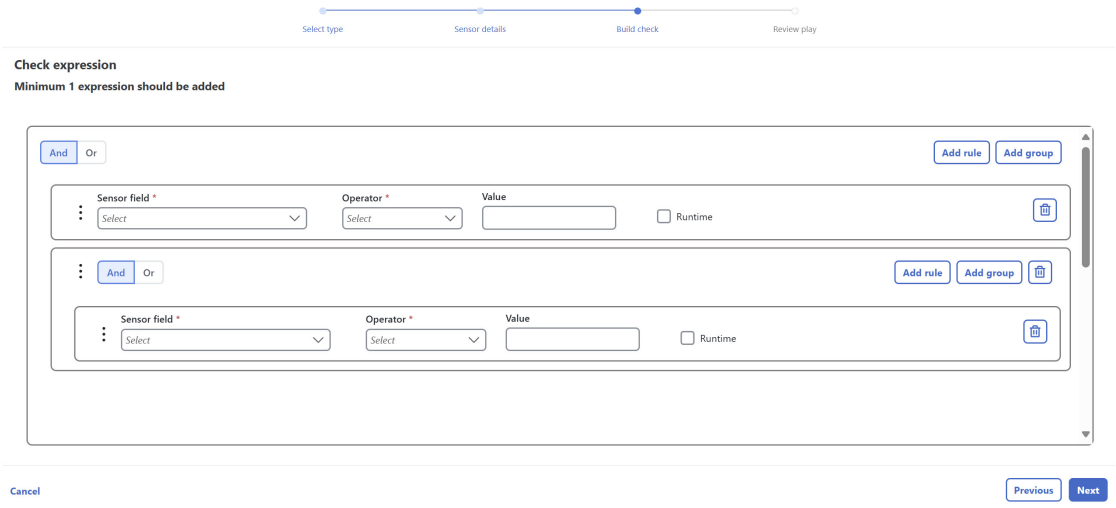

Click **Next** to continue.

Step 4 **Creating a Device config play or Service play**

Ensure that the configuration you are trying to create is available in NSO; otherwise, it will show an error.

When creating a Service play, you are not creating a new service for NSO but creating a play to manage and provision an existing service in one or more NSO instances. For more information, see https://developer.cisco.com/docs/nso/.

a) In the **Configure play** window, click or the **Import** link to import your device config (.JSON) file. You can download and use the sample configuration template. Browse and select your .JSON file, and click **Import**.
b) In the acknowledgment prompt, click **Continue** to select the NSO instance for the config you have imported.
c) Select the NSO provider instance from the dialog box and click **Process Payload**.

Figure 10: Select NSO provider

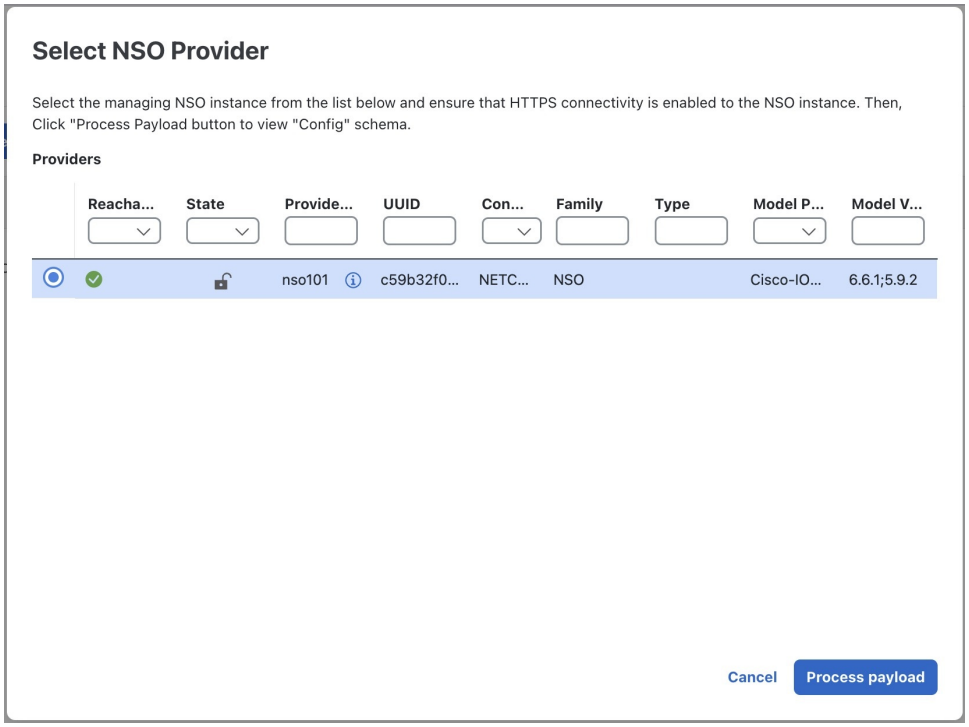

Note
The creation workflow of a Service play is similar to the Device config play, except in the template of the payload file used.

d) The **Configure play** window opens, displaying information from the payload file. You can edit the *value* or *description* columns with the values that you want to see during a playbook execution.

Figure 11: Configure play

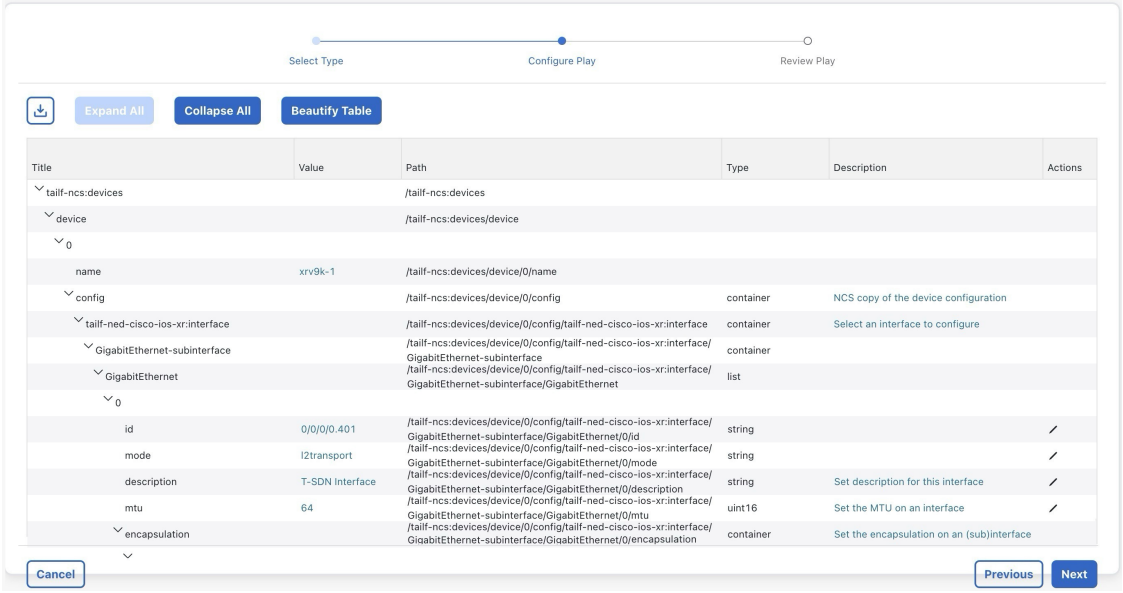

Click **Next** to continue.

Step 5 In the **Review play** window, review the parameters of your play. Click **Dry run** to validate your parameters.

Label your play with a unique **Name** and **Description**.

Note
Cisco also formats the play names with indicators such as cfg for configuration, chk for check, and so on, in the name to help you organize the plays properly. You can also use similar tagging for the plays you create.

You can also add labels to your play to group it in the future (optional).

Note
The labels determine the type of devices with which you can use the play. For example, an IOS XR play cannot run on IOS XE devices. Be sure to review the labels (IOS XR, IOS XE, and so on) when you add them.

Figure 12: Review play

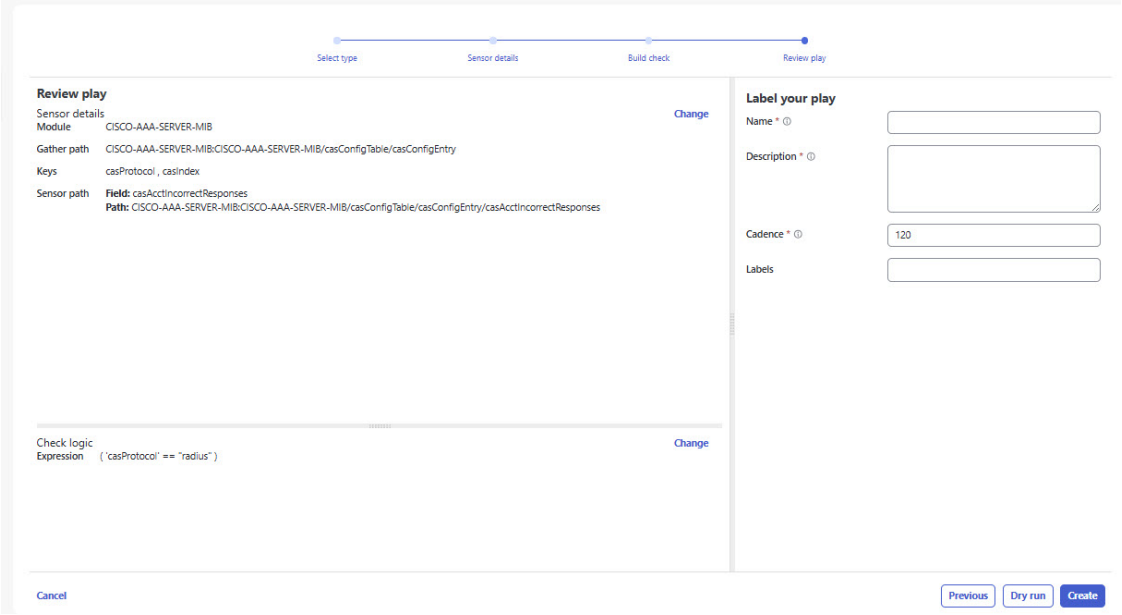

Step 6 If you are satisfied with your changes, click **Create**.

The **Play List** window opens, displaying your new custom play in the play list.

Export plays

A user must have Change Automation read permission to export any custom play authored or imported by you or another user into Change Automation.

The exported archive contains only the user-customizable files listed in Playbook components and files, on page 28. Once you extract them from the archive, you can identify the play components by their file names and filename extensions.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Play List**.

Step 2 Check the check boxes for the custom plays you want to export.

Step 3 Click . Your browser will prompt you to select a path and the file name when saving the gzipped tar archive. Follow the prompts to save the file.

Import custom plays

You can import any custom play that meets the following requirements:

- The play files must be packaged as a gzipped tar archive.
- The archive must contain a `.play` file (a data spec file for the play), at minimum.
- The archive file must have a unique name.

Note For more details about editing and importing, see Cisco Crosswork Change Automation Developer Guide.

You can overwrite a custom play. The system will warn you when you are about to overwrite a custom play but will not prevent you from doing so.

Warning Take precautions to ensure you do not accidentally overwrite the custom plays you created.

Before you begin

To import plays, a user must have write access. For more information about granting a user read-write role access, see Verify installation and configure system settings, on page 5.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Play List**.

Step 2 Click . Your browser will prompt you to browse to and select the gzipped archive file containing the plays you want to import.

Make sure that there are no Cisco-supplied plays with the same name as the play you intend to import. If you import a play with the same name, it will fail.

Step 3 Follow the prompts to import the archive file.

Delete custom plays

You can delete custom plays only. You cannot delete a Cisco-supplied play.

Your user ID must have Change Automation delete permission to delete plays.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Play List**.

Step 2 In the **Play List** window, select the custom plays you want to delete.

Step 3 Click the icon.

Step 4 When prompted, click **Delete** again to confirm.

About customizing playbooks

You can create your own playbooks from scratch, based on details from Cisco-supplied playbooks. You can also create custom playbooks using the available plays.

Creating and modifying Cisco-supplied playbooks are engineering tasks that take place outside of the user interface for Change Automation. As such, they are outside the scope of this User Guide.

Cisco supplies developer-level documentation for Cisco-supplied playbooks. For more information on how to create custom plays and playbooks, see Developer Guides on Cisco DevNet.

Playbook components and files

Change Automation playbooks contain various components, referred to using specialized names. The components are implemented in the playbook as files. Some of these components' names are borrowed from the Ansible specification, but all have their definitions, and not all the corresponding files can be customized by users. Some components are Cisco proprietary intellectual property; while you can use them in custom plays and playbooks, you cannot customize them directly. For more information, see Writing custom playbooks at Cisco Crosswork Change Automation developer guide.

Create a custom playbook through the UI

Change Automation allows users with admin and read/write roles to create custom playbooks using the available plays. For more information about granting read/write role access to a user, see the section "Assign Change Automation User Access Levels" in the topic Verify installation and configure system settings, on page 5.

Note You cannot edit a custom playbook once it is created. We recommend that you perform a dry run of the playbook *before completing the creation* to ensure that the playbook's purpose is met. Once created, if you must make changes to the custom playbook, you have to recreate your playbook with the relevant changes.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Playbook List**. The **Playbook List** window is displayed.

Step 2 Click  to create a custom playbook. The **Select plays** window opens displaying the available plays.

Figure 13: Select plays

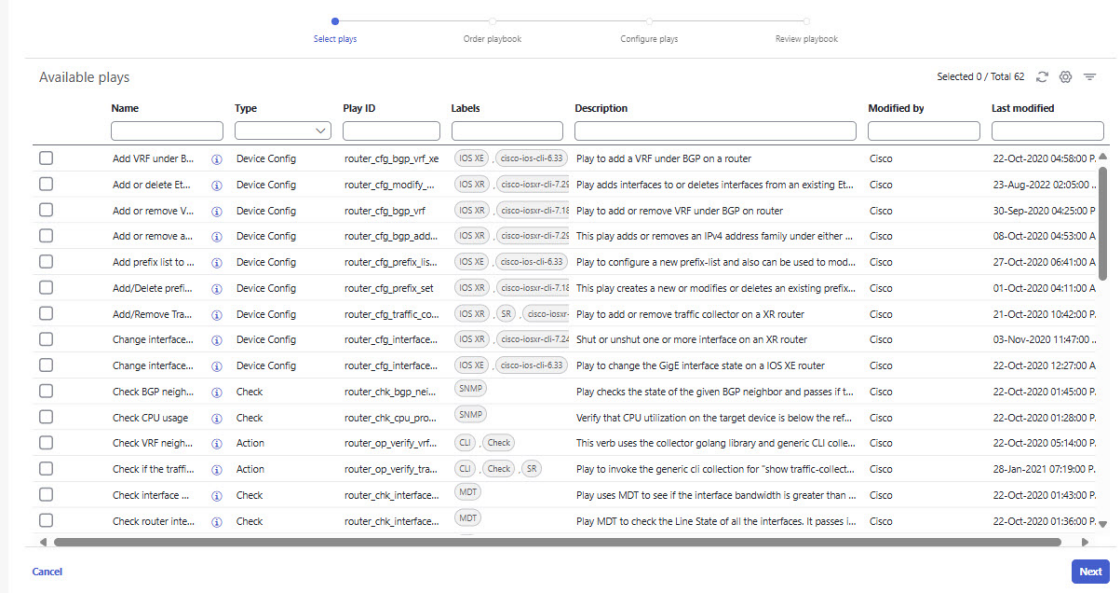

Select all the plays you want in your playbook, and click **Next**.

Note
The recommendation is to include the **Perform Check Sync on the device** and **Sync NSO from device** plays as a pre-step to running other operations in the playbook or as part of pre-maintenance.

The **Perform Check Sync on the device** play checks the device sync status with NSO and performs sync only when needed, based on the playbook's sync parameter value. It reduces the playbook execution time and ensures the NSO configuration matches the device configuration.

- If the playbook's sync parameter is set to True and the device is not in sync, the **Perform Check Sync on the device** play will sync the device with the NSO configuration.
- If the sync parameter is set to False and the device is not in sync, the play will fail to execute with a commit message.
- If the device is already in sync, the play will succeed.

Step 3 In the **Order playbook** window, arrange the order of the plays in the playbook as per the execution phase (Continuous, Pre-Maintenance, Maintenance, Post-Maintenance). By default, all the selected plays are displayed within the Maintenance phase. You can click and drag the plays to rearrange them to the appropriate phase.

Depending on the type of play you have selected, it may be restricted from being used in certain phases. For example, a configuration play cannot be used outside of the Maintenance phase.

For more information on each execution phase, see Playbook execution order, on page 37.

Change Automation also formats the play names with indicators such as "cfg" for configuration, "chk" for check, and so on, in the name to help you organize the plays properly. You can use similar tagging for the plays you create.

You can also duplicate or delete a play by clicking on the icons provided.

Figure 14: Order playbook

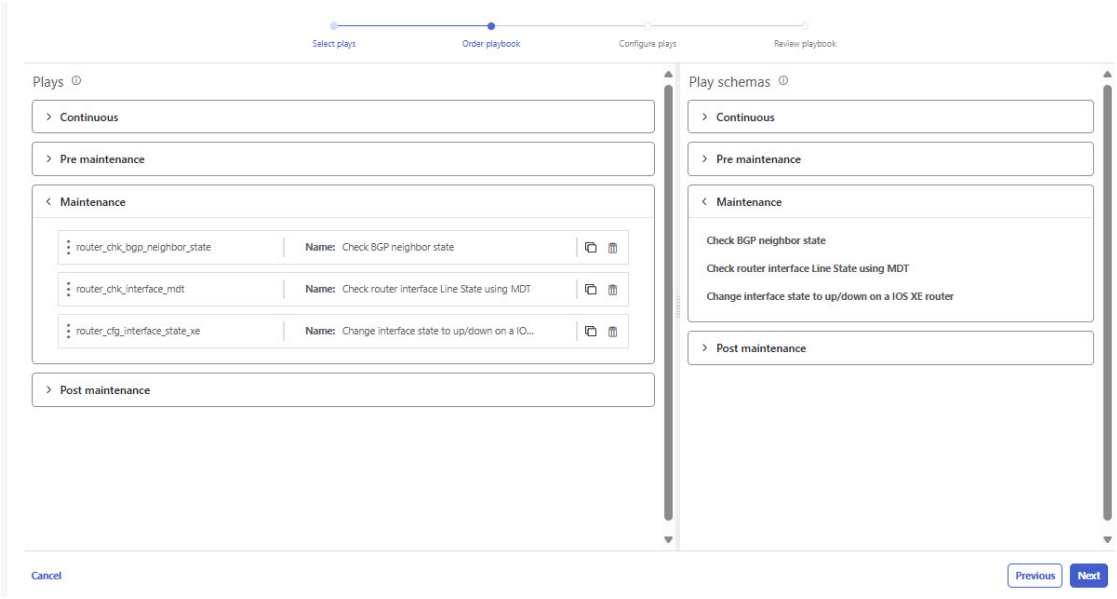

Click **Next**.

Step 4 The **Configure plays** window opens, displaying the plays in each execution phase and the play schemas.

Figure 15: Configure plays

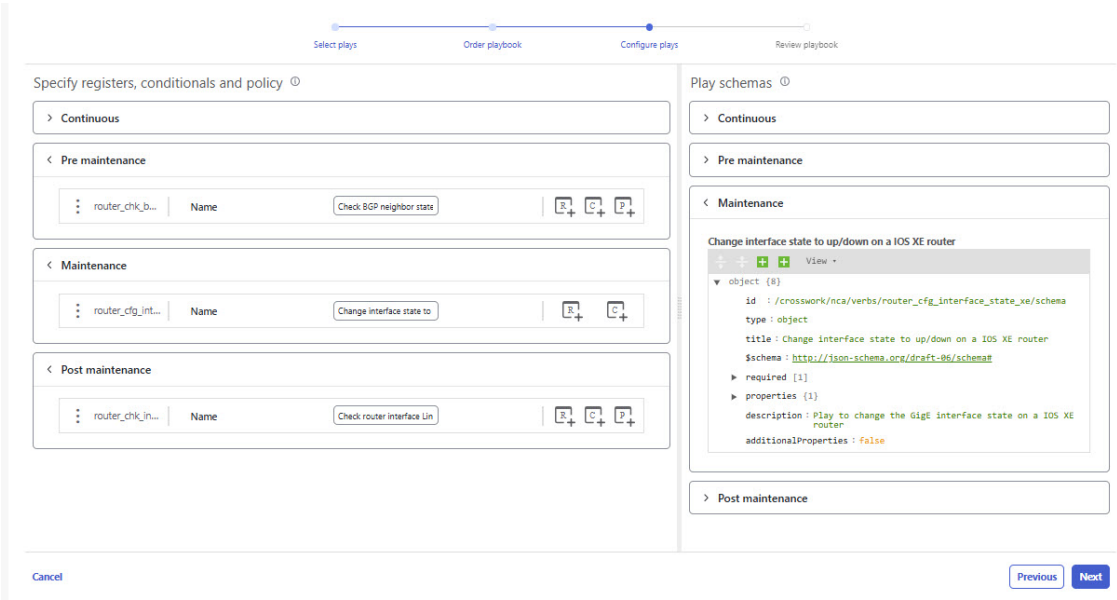

You can perform the following:

- Click to specify a policy for a play. In the **Specify policy** dialog box, specify relevant values for the fields provided. Click for more information about each field. Click **Save** to save your policy values.

  Note

Policies are applicable to Check plays.

Figure 16: Specify policy

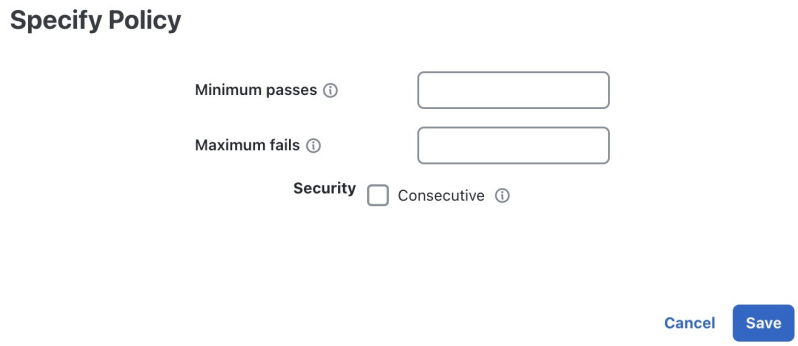

- Click to apply a condition to a play. Execution of the play proceeds only if the condition is met. In the **Specify Conditionals** dialog box, click **Add Condition** to add a condition. Click **Save** to save your conditional values.

Figure 17: Specify conditionals

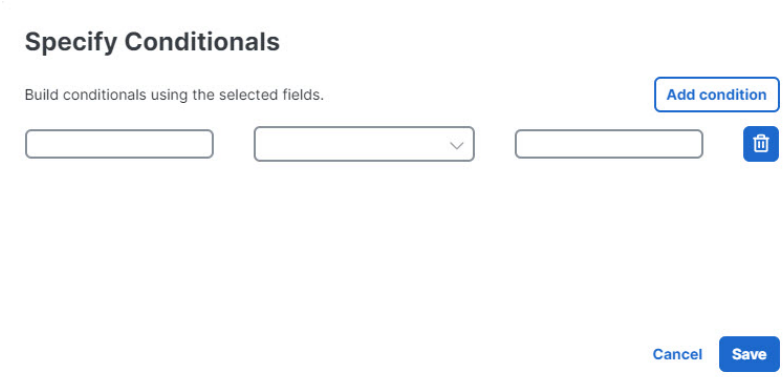

- Click to specify a register for a play. Specifying registers allows you to use the output of a previous play as the input for another play. Click **Save** to save your registers.

Figure 18: Specify registers

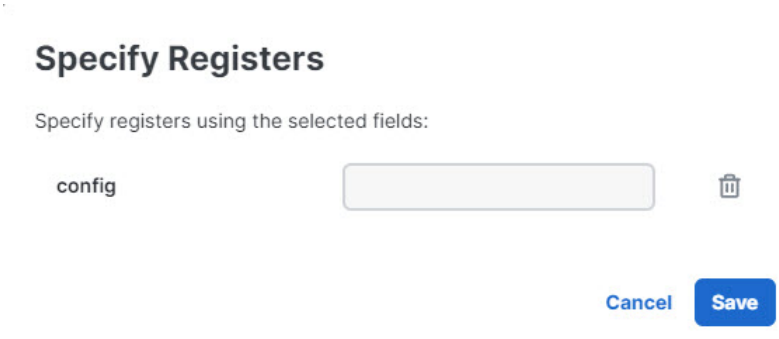

- (Optional) Rename the plays if you want them to be displayed with different names during the playbook execution.

Click **Next** to continue.

Step 5 In the **Review playbook** window, review the plays in your playbook. Enter relevant values for the **Playbook details** fields. You can click ⓘ for more information about each field.

Figure 19: Review playbook

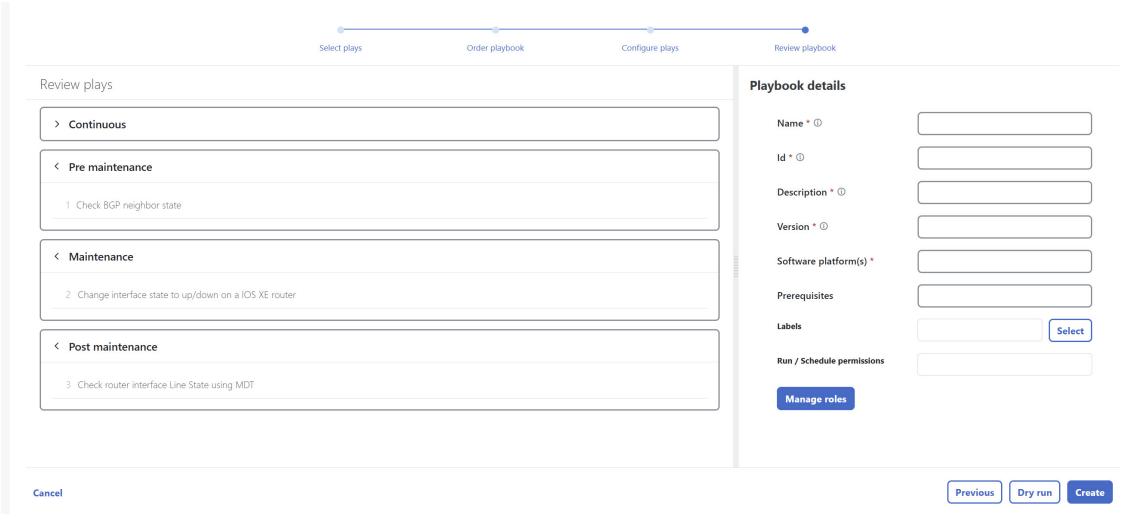

Note
For the **Software platform(s)** field, make sure to use the exact software type name as it is mentioned in **Device Management** > **Network Devices** > **Software type** column.

Step 6 (Optional) Click **Select** and perform one of the following, as applicable, to set the **Labels**:

- Select the applicable label and click **Done**.
- Click **+ New label**, enter relevant values for **Label** and **Roles**, and click **Save**. Select the new label and click **Done**.

Note
The labels determine which user or roles can run which playbooks.

For more information on assigning playbooks to specific roles, see Assign playbooks to specific roles, on page 36.

Step 7 (Optional, but recommended for testing the playbook) After you enter the relevant details, click **Dry run** to validate the parameters. A dialog box opens, displaying the playbook Details.

Figure 20: Playbook details

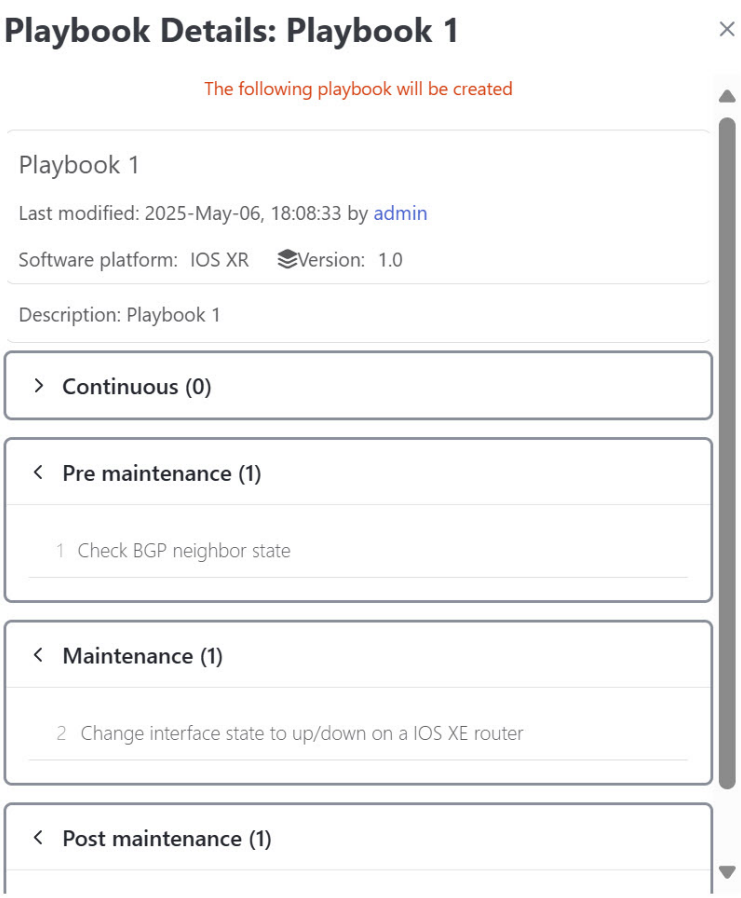

Note
Dry run does not commit the changes but provides a platform to validate whether the playbook would work with the parameters you entered.

Step 8 Click **Previous** to navigate back to a step to make changes as necessary to get the playbook to function properly.

Step 9 Click **Create** to create the playbook.

The **Playbook list** window opens, displaying your new custom playbook on the list.

Create a custom playbook using APIs

This section explains the steps to create a custom playbook using APIs.

Note A playbook containing a custom play can be created through the UI (see Create a custom playbook through the UI, on page 28) or using APIs.

A playbook consisting of one or more custom plays is expected to have a *dataspec* value for the custom play in the playbook file. The *dataspec* value is generated when the custom playbook is created using the API in this procedure. You cannot create the same custom playbook using the import option (API: **/v1/mops/import**), as it does not add the *dataspec* value for the custom play.

Procedure

Step 1 Ensure that the plays (stock or custom) you need for the playbook are created beforehand.

You can create a custom play either through the UI (see Create a custom play using templates, on page 21) or using API (use the API call **//crosswork_ip:30603/crosswork/nca/v1/Plays/device/config**).

Note
If you are importing plays that share the same name with existing plays, then the error "`Play validation failed, custom Play already present`" will be displayed, to prevent the existing plays being overwritten.

Step 2 Create the playbook using the following API:

API call: **//crosswork_ip:30603/crosswork/nca/v1/mops**

Export playbooks

You can export any playbook as a gzipped tar archive. This includes any Cisco-supplied playbook and custom playbooks you or another party have authored and imported into Change Automation.

The exported archive contains only the user-customizable files listed in Playbook components and files, on page 28. It also contains one or more `.pb` files (for example, `router_config_bgp_rd.pb` for the playbook code), which are parsed and processed at the back end.

You can edit the exported files as needed. Then, you can import them as explained in Import playbooks, on page 35.

Your user ID must have Change Automation read permission to export playbooks and write permissions to import new or modified playbooks.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Playbook List**.

Step 2 (Optional): In the **Playbook List** window, filter the table as needed.

Step 3 Check the check boxes for the playbooks you want to export. Check the check box at the top of the column to select all playbooks for export.

Step 4 Click . Your browser will prompt you to select a path and the file name when saving the gzipped tar archive. Follow the prompts to save the file.

Import playbooks

You can import any custom playbook, provided it meets the following requirements:

- The playbook files must be packaged as a gzipped tar archive.
- The archive must contain a `.pb` file, at minimum.
- The archive file must have a unique name.

The individual files included in the archive must meet the additional validation requirements. See Import a custom Mop (playbook) on Cisco DevNet.

Note While you cannot overwrite a Cisco-supplied playbook, you *can* overwrite a custom playbook. The system will warn you when you are about to overwrite a custom playbook but will not prevent you from doing so. Take precautions to ensure that you do not overwrite your custom playbooks accidentally.

You cannot re-import an exported Cisco-supplied playbook with the same name as the original.

Before you begin

To import playbooks, a user must have write access. For more information about granting a user read-writer role access, see Verify installation and configure system settings, on page 5.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Playbook List**.

Step 2 Click . Your browser will prompt you to browse to and select the gzipped archive file containing the playbooks you want to import.

Make sure there is no existing playbook with the same name as the playbook you intend to import unless you want to overwrite the existing playbook.

If you are creating an improved version of a playbook, it is recommended that you use a version number or other indicator to ensure that the name is unique and does not overwrite the original playbook until the replacement is completely tested.

Step 3 Follow the prompts to import the archive file.

Delete custom playbooks

You can delete user-defined playbooks only. You cannot delete a Cisco-supplied playbook.

Your user ID must have Change Automation delete permission to delete playbooks.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Playbooks List**.

Step 2 In the **Playbooks List** window, select the custom playbook you want to delete.

Step 3 Click the icon.

Step 4 When prompted, click **Delete** again to confirm.

Assign playbooks to specific roles

This section explains how to assign playbook labels to specific roles so that they can run and import the playbooks with that particular label. Admin users can enable other users to run playbooks with a specific label.

Before you begin

If required, create a new user to whom you would want to assign the playbook. For more information, see the topic Create user roles in the *Cisco Crosswork Network Controller Administration* guide.

Procedure

Step 1 Go to **Administration** > **Users and Roles** > **Roles** > **Playbook**.

Step 2 Under the **Roles** pane, select the role to whom you want to assign the playbook labels.

Step 3 Enable the **Permissions** check boxes for the **Playbook label**(s) you want to assign.

Figure 21: Playbook labels

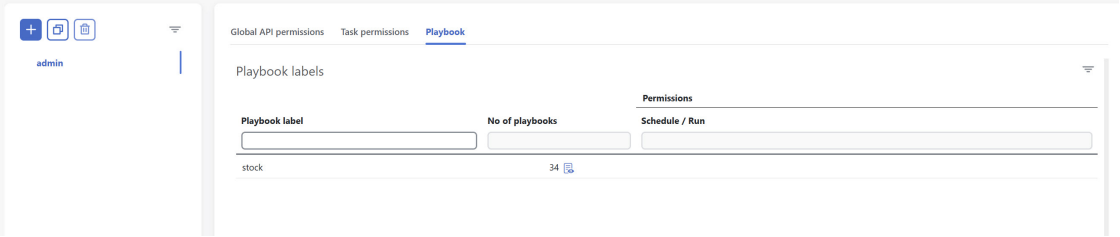

About running playbooks

You must have permission to run the playbooks with a particular playbook label. For more information on assigning playbooks to specific roles, see Assign playbooks to specific roles, on page 36.

Running any playbook consists of five steps:

1. Select the **playbook** you want to run (see View the playbook list, on page 19).

2. Select the **device or devices** that you want to run it on.
3. Enter the appropriate run-time **parameters** that you want the playbook to apply.
4. Select the **execution mode** that you want to use:
   a. Perform a dry run of a playbook, on page 38, where you can see what the playbook does before make changes to the network.
   b. Run playbooks in single stepping mode, on page 44, so you can pause after each playbook check or action, and roll back changes you did not intend.
   c. Run playbooks in continuous mode, on page 50 and apply the changes immediately.

   While selecting the execution mode, you can also choose to:
   - Schedule playbook runs, on page 56 for another calendar date or time.
   - **Collect syslogs** during and after the run. Syslog collection is available only when running the playbook in single-stepping or continuous execution mode and only if you have already configured a syslog storage provider.
   - Specify a **Failure policy**, where you decide what the system should do if a failure occurs during the playbook run.
5. **Confirm** your settings and run the playbook in the execution mode you selected.

Depending on their complexity and network factors, some playbooks may take much time to run. You can view the run details and status at any time during and after the completion of a run. If the playbook is still running, you can also choose to cancel it. For details, see View or abort playbook jobs, on page 58.

Playbook execution order

When it is running, every playbook conducts checks and configuration changes in four phases, which correspond to sections of the playbook code (identified using the tags discussed in Playbook components and files, on page 28).

1. **Pre-Maintenance**—This phase of the playbook includes non-disruptive checks and any other operations on the device that prepare it for potentially traffic-impacting changes. For example:
   - Take snapshots of various routing protocol states.
   - Take snapshots of memory, CPU, and system health parameters.
   - Validate the capacity (storage, memory) on active and standby routers for the new software patch upgrade.
2. **Maintenance**—This phase of the playbook includes any task that may disrupt traffic flowing through the router or impact neighboring routers. For example:
   - Cost out the router and wait until traffic drains out completely.
   - Verify that the redundant router is healthy and carrying traffic.
   - Perform the upgrade procedure on the device.
   - Reconfigure the device(s) to support a new configuration or feature.

3. **Post-Maintenance**—This phase of the playbook includes verification tasks to perform on the router after any disruptive operation. For example:
   - Verify that the current state matches the desired state.
   - Cost in the router and wait for traffic to return to normal levels.
4. **Continuous**—In addition to the three serial phases already described, Change Automation can also run check tasks that span the entire duration of playbook execution. These tasks check the state of the router while the playbook is being deployed and cancel the playbook execution if any catastrophic or undesirable state change occurs. The checks in the playbook may also monitor a neighboring router to guarantee no second-order failures in the network while the changes are being deployed.

Perform a dry run of a playbook

A dry run lets you view configuration changes that the playbook will send to the device without performing the actual commit of the changes, as you would with a run in the single-stepping or continuous execution modes.

It is a best practice to perform a dry run and verify the configuration changes before you deploy those changes to the router. If the dry run fails, you may want to debug its parameter values using another dry run. You can also debug by performing a single-stepping run, which will allow you to abort and rollback changes after one or more of the plays, instead of only at the end, as part of a continuous run's Failure Policy.

Note that dry run mode is intended for use only with playbooks that perform device configuration changes via Cisco NSO. Some playbooks do not support dry run mode. For example, **Install a SMU or an optional package on a router** and **Uninstall an optional package or a SMU**.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Run Playbook**.

Step 2 In the **Available playbooks** list on the left, click on the playbook you want to dry run. On the right, the window displays the playbook name, hardware and software compatibility information, and descriptions for all the plays in the selected playbook.

Figure 22: Select playbook

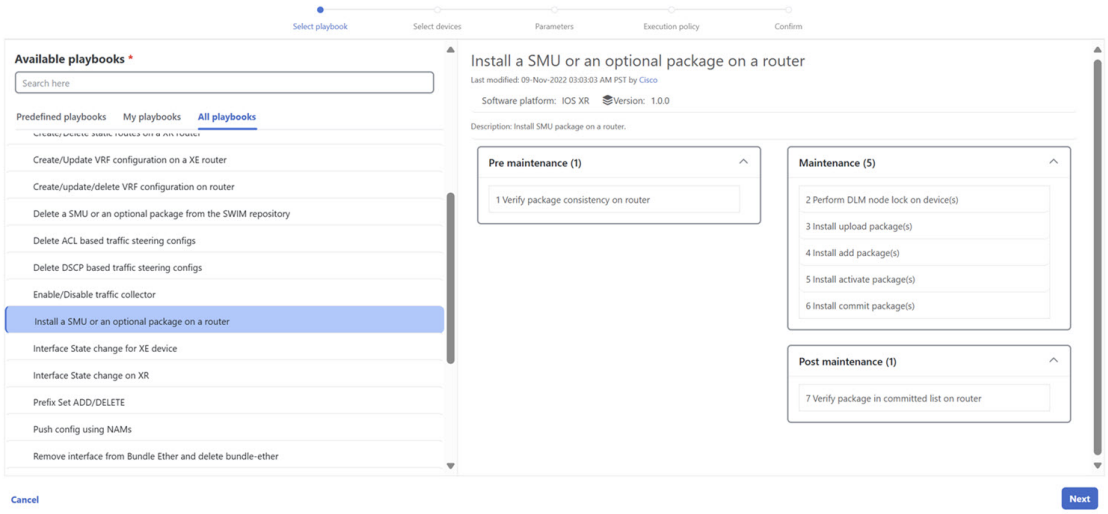

Step 3 Click **Next**. The **Select devices** window appears.

Figure 23: Select devices

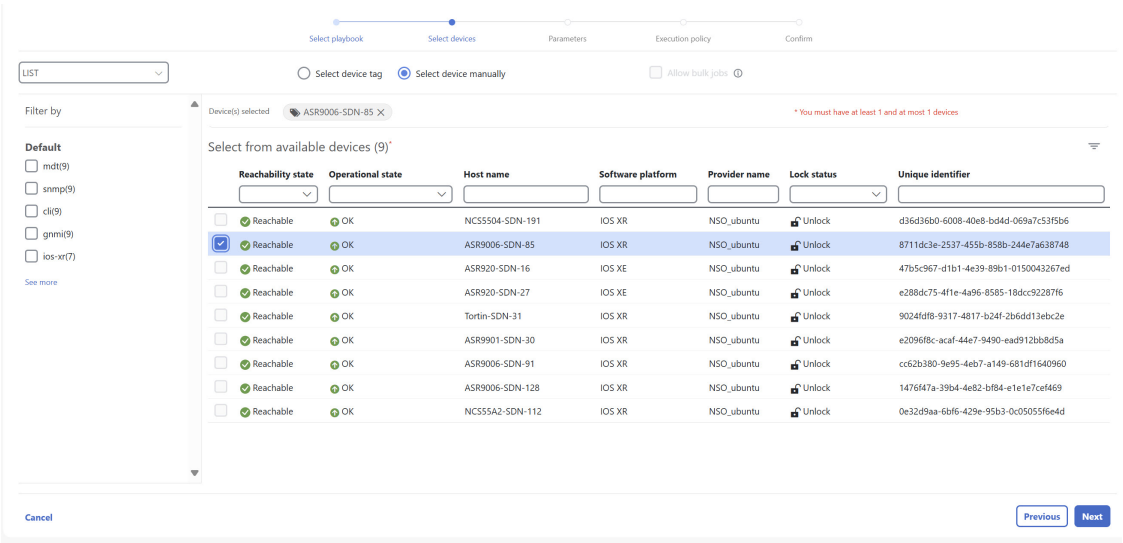

Using this window:

- You can toggle between the table view and topology map view by clicking and selecting the relevant option on the drop-down button on the upper left corner of the window. Choose **Select devices from list** or **Select devices from map** to select the table view or topology map view respectively. By default, the table view is displayed.
- With the topology map view displayed, you can toggle between the map's geographical and logical views by clicking on the or the . You can also zoom, display bandwidth utilization, and change logical view layouts as you do with the topology map you see when you select **Topology** from the main menu.
- You can select the devices manually or using tags. The **Select device tag** option targets you to select the relevant tag instead of devices from the table and all devices associated with the relevant tag are selected. The **Select device**

manually option allows you to select the devices from the list using quick and advanced filters and filter by tags on the left. Hover the mouse pointer over the ⓘ icon next to the options for more information. You can also view the selection criteria, such as the number of devices required for the selected playbook.

Note
If you are a non-admin user and selecting the devices manually, make a note of the following:

- The devices on which you want to run the playbook must belong to a Device Access Group, and you must have access to this Device Access Group. For more information on Device Access Groups and associating a user with a Device Access Group, see the Manage Device Access Groups section in the *Cisco Crosswork Network Controller Administration* guide.
- If your role is associated with an empty Device Access Group, then you will receive an error message.
- If your role is associated with multiple Device Access Groups and the device belongs to any of these Device Access Groups, then you can run the playbook on this device. If the device does not belong to any of your Device Access Groups, then the operation fails.
- If you are selecting multiple devices (using **Allow bulk jobs** option or using tags) and if any of the devices does not have access, then an error message appears stating that this list of devices does not have access to run the playbook.

- In the **Select device manually** selection mode, you can check the **Allow bulk jobs** check box to select multiple devices and run the selected playbook on them simultaneously. Based on your selection, the system creates a static group of multiple jobs. Hover the mouse pointer over the ⓘ icon next to the check box for more information. There is no limit to the number of devices you can select for a bulk job.

Note
Allow bulk jobs option is enabled for playbooks that can be executed on a single device.

Step 4 Click **Next**. The **Parameters** window appears.

Step 5 In the fields provided in the **Parameters** window, enter the playbook parameter values for this dry run.

Figure 24: Parameters

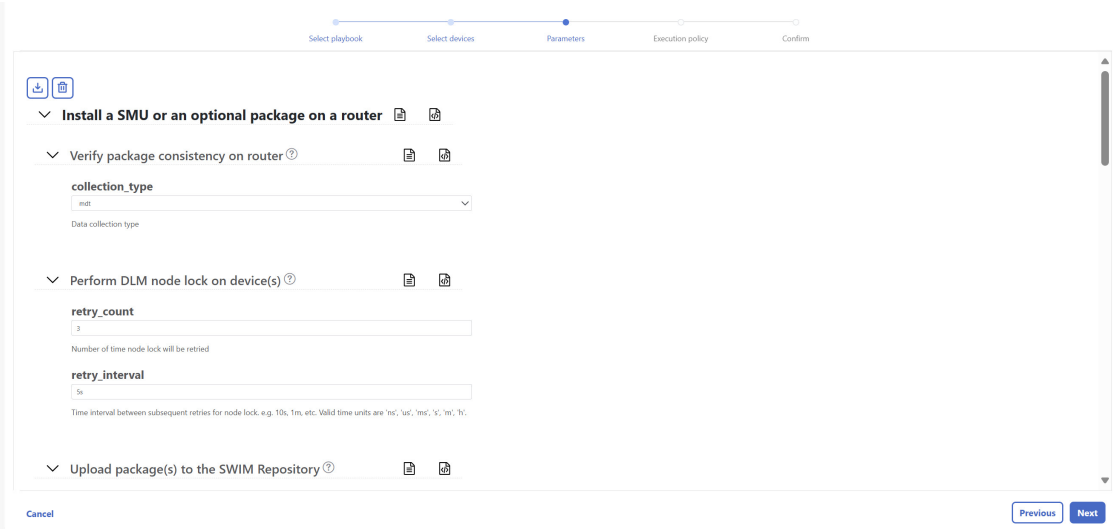

With the **Parameters** window displayed, you can also:

- Click to upload a JSON file with the parameter values you want. You will be prompted to navigate to the JSON parameters file you have previously prepared (or downloaded from a previous playbook run) and then upload it as appropriate for your browser and operating system.
- Click to enter the parameter values in JSON format. A popup text window displays the full list of JSON parameters with empty values in quotes. Edit the values, and when you are finished, click **Save**.

Figure 25: Edit JSON

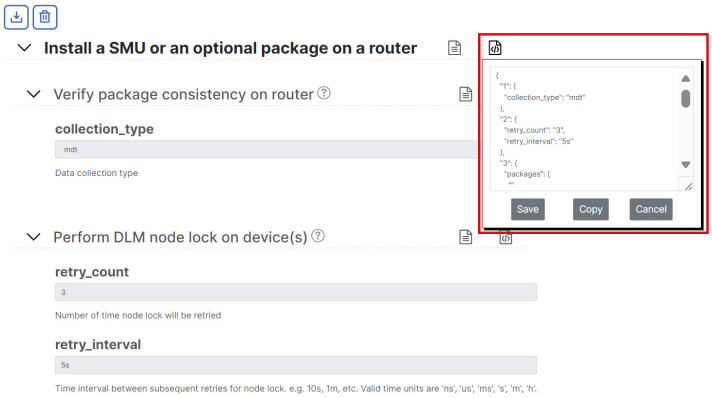

- Click to select the items you want to configure. A pop-up text window will appear, displaying the complete list of plays or parameters that can be configured. For the playbook, it will show a list of plays and, for each selected play, it will display a list of parameters that can be configured for that play. If you unselect any item, it will not appear for configuration. Uneditable options indicate mandatory items that cannot be deselected.

Figure 26: Object properties for a playbook

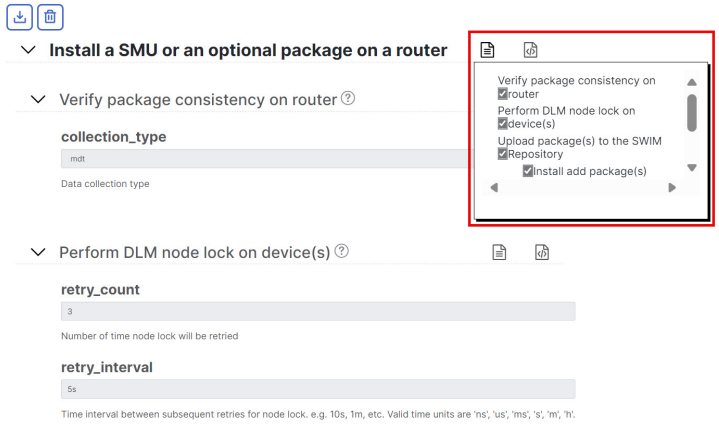

Figure 27: Object properties for a play

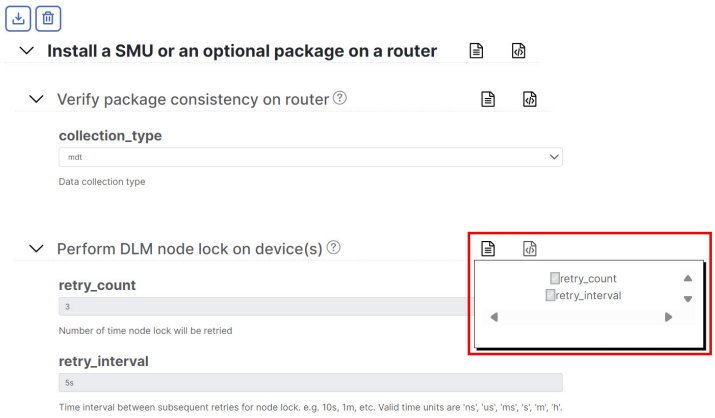

- Click ⊕ to add additional instances of a particular parameter, if required for the playbook you are running. Click 🗑 to delete instances added in this way.
- Click 🗑 to clear all the parameter values entered so far.

Step 6 With the parameter values set, click **Next**. The **Execution policy** window appears.

Figure 28: Execution Policy

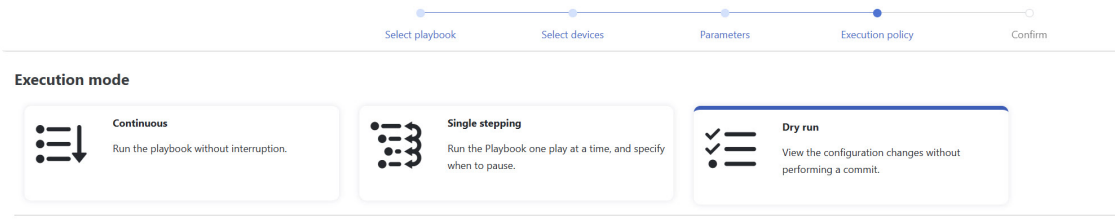

Step 7 Choose **Dry run** and click **Next**. The **Review your job** window appears, displaying a summary of your choices: playbook, devices, parameters, and execution policy.

Figure 29: Review your Job

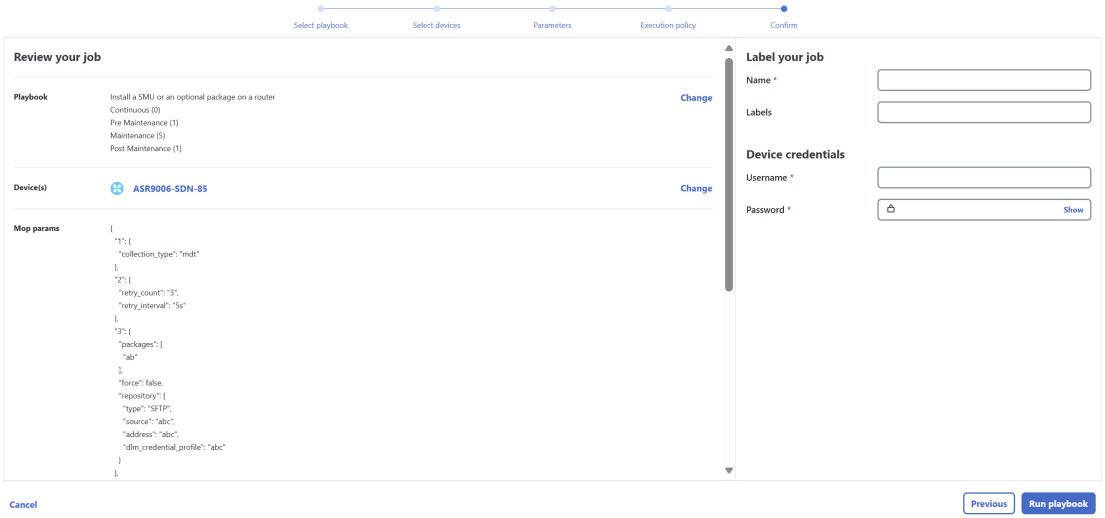

In this window:

- You must provide a relevant **Name** for the job.
- You can enter labels for your job using the **Labels** field.
- You can click on any **Change** links in the **Review your job** window summary to modify your choices.

Step 8 (Optional) Enter the device credentials (name and password).

Note
This step is applicable only if **Credential prompt** is enabled in the Change Automation settings. For more information, see Verify installation and configure system settings, on page 5.

Step 9 When you are ready to continue, click **Run playbook**.

Step 10 At the confirmation prompt, click **Confirm**. The **Execution mode** window is displayed.

Step 11 After the dry run is complete:

- Click the **Dry run** tab and verify the configuration changes that would be pushed to the device had this not been a dry run. This tab will display a `no config change` message if no changes have been made. Please note that this tab shows only cumulative configuration changes, not each individual change made. For example, if a playbook configures `set-overload-bit` in one step and then unconfigures it using `no set-overload-bit` later, the tab will show `no config change`.
- Click the **Events** tab to see success and failure messages for each step of the playbook. This can help you diagnose and correct problems with individual plays and the run as a whole. For troubleshooting information, see Troubleshoot change automation, on page 61.
- Click the **Console** tab to see messages that are generated during the run.

As syslog collection is disabled for dry run, the **Syslog** tab will contain only a message stating that.

Step 12 (Optional) If you want to perform a single-step debugging run, or are ready to commit the changes to the device, click **Execute now**. The **Execution policy** window will display all of your parameter values from the dry run.

Run playbooks in single stepping mode

Single-stepping execution mode is a handy way to test a custom or modified playbook or diagnose problems with a pre-packaged playbook that does not give you the desired results. Unlike a dry run, a single-stepping execution commits configuration changes to the device as the playbook runs. However, you can set breakpoints on or pauses after any Maintenance or Post-Maintenance action in the playbook. Note that while you can set breakpoints on Pre-Maintenance actions, doing so will have no effect, and these actions will not pause.

Whenever the playbook hits a breakpoint, it will stop and will not continue until you issue the command to proceed. At each pause, you can also abort the entire run and roll back all changes made or roll back to any previous play.

Procedure

Step 1 From the main menu, choose **Network Automation** > **Run Playbook**.

Step 2 In the **Available playbooks** list on the left, click on the playbook you want to run. On the right, the window displays the playbook name, hardware and software compatibility information, and descriptions for all the plays in the selected playbook.

Figure 30: Select playbook

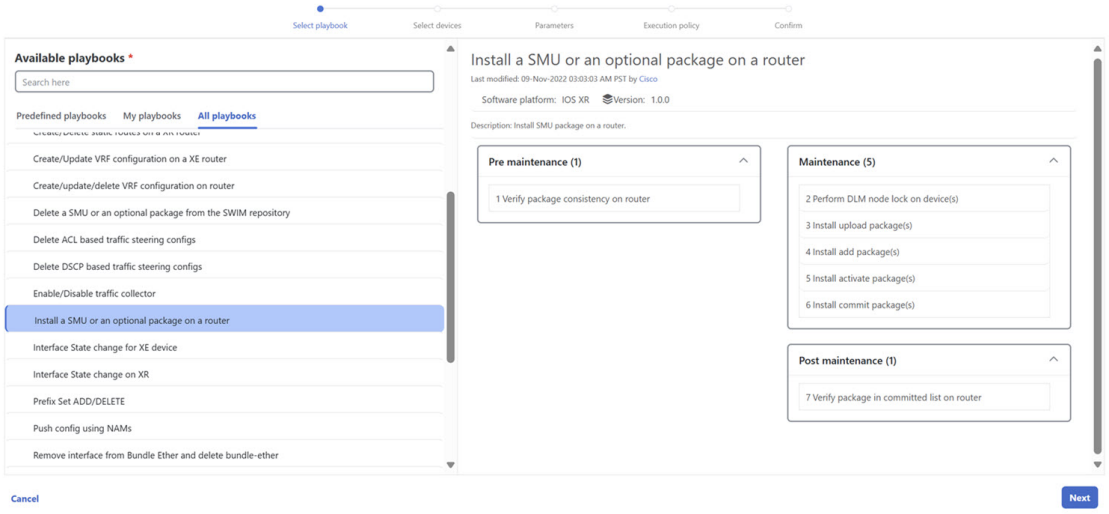

Step 3 Click **Next**. The **Select devices** window appears.

Figure 31: Select devices

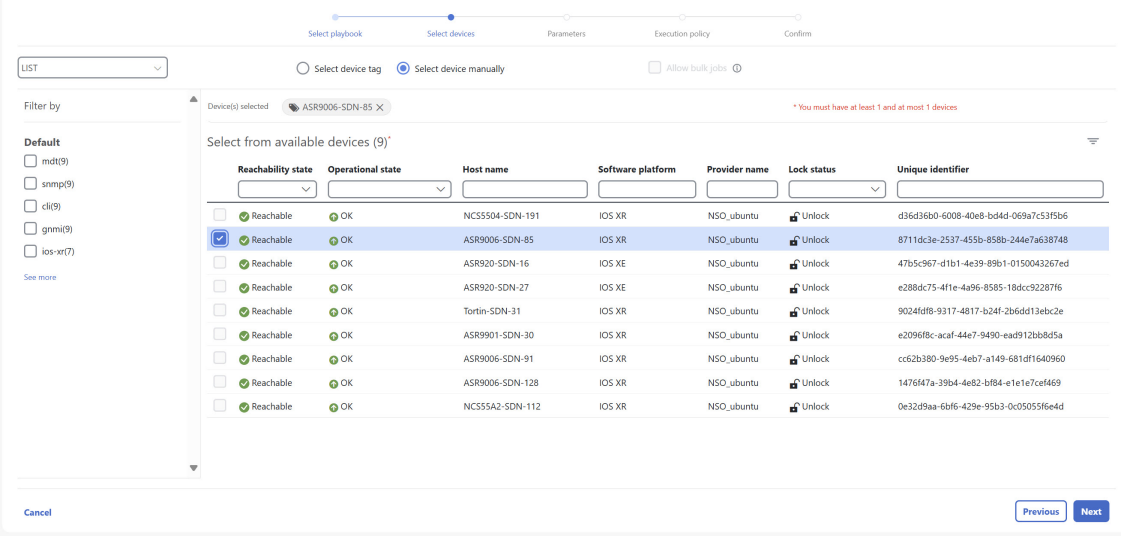

Using this window:

- You can toggle between the table view and topology map view by clicking and selecting the relevant option on the drop-down button in the upper left corner of the window. Choose **Select devices from list** or **Select devices from map** to select the table view or topology map view, respectively. By default, the table view is displayed.
- With the topology map view displayed, you can toggle between the map's geographical and logical views by clicking on the or the . You can also zoom, display bandwidth utilization, and change logical view layouts as you do with the topology map you see when you select **Topology** from the main menu.
- You can select the devices manually or using tags. The **Select device tag** option targets you to select the relevant tag instead of devices from the table and all devices associated with the relevant tag are selected. The **Select device manually** option allows you to select the devices from the list using quick and advanced filters and filter by tags on the left. Hover the mouse pointer over the icon next to the options for more information. You can also view the selection criteria, such as the number of devices required for the selected playbook.

  Note
  If you are a non-admin user and selecting the devices manually, make a note of the following:

  - The devices on which you want to run the playbook must belong to a Device Access Group, and you must have access to this Device Access Group. For more information on Device Access Groups and associating a user with a Device Access Group, see the Manage Device Access Groups section in the *Cisco Crosswork Network Controller Administration* guide.
  - If your role is associated with an empty Device Access Group, you will receive an error message.
  - If your role is associated with multiple Device Access Groups and the device belongs to any of these Device Access Groups, then you can run the playbook on this device. The operation fails if the device does not belong to any of your Device Access Groups.
  - If you select multiple devices (using the **Allow bulk jobs** option or using tags) and if any of them does not have access, an error message appears stating that this list of devices does not have access to run the playbook.

- In the **Select device manually** selection mode, you can check the **Allow bulk jobs** check box to select multiple devices and run the selected playbook on them simultaneously. Based on your selection, the system creates a static group of multiple jobs. Hover the mouse pointer over the ⓘ icon next to the check box for more information. There is no limit to the number of devices you can select for a bulk job.

  Note
  Allow bulk jobs option is enabled for playbooks that can be executed on a single device.

Step 4 Click **Next**. The **Parameters** window appears.

Step 5 In the fields provided in the **Parameters** window, enter the playbook parameter values for this run.

Figure 32: Parameters

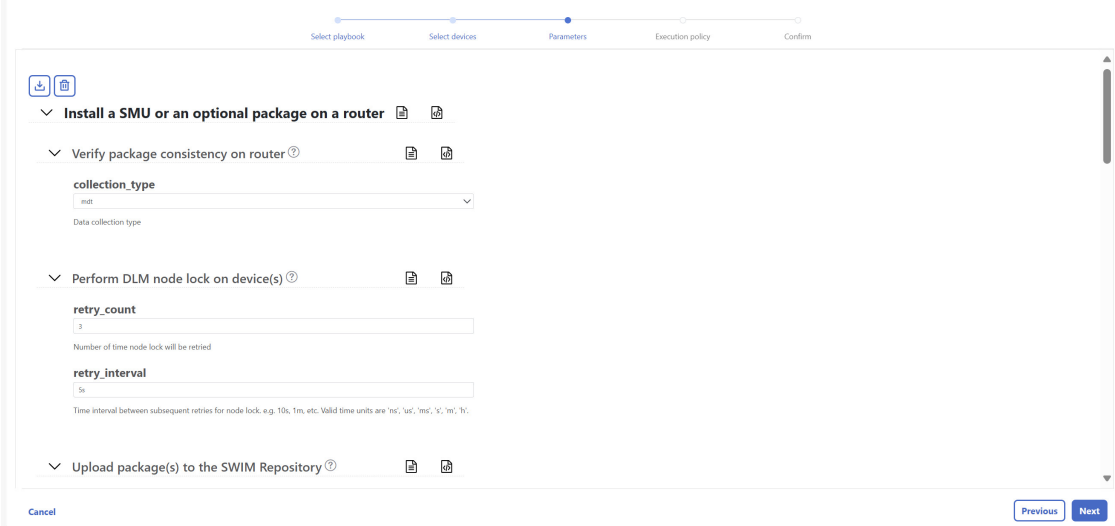

With the **Parameters** window displayed, you can also:

- Click ⬇ to upload a JSON file with the parameter values you want. You will be prompted to navigate to the JSON parameters file you have previously prepared (or downloaded from a previous playbook run) and then upload it as appropriate for your browser and operating system.

- Click </> to enter the parameter values in JSON format. A popup text window displays the full list of JSON parameters with empty values in quotes. Edit the values, and when you are finished, click **Save**.

Figure 33: Edit JSON

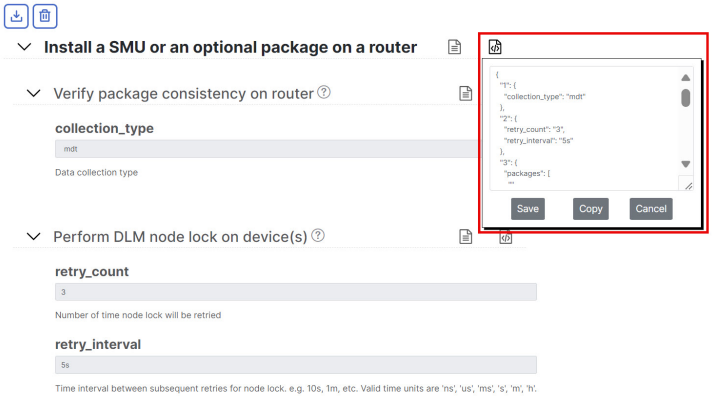

- Click  to select the items you want to configure. A pop-up text window will appear, displaying the complete list of plays or parameters that can be configured. For the playbook, it will show a list of plays and, for each selected play, it will display a list of parameters that can be configured for that play. If you unselect any item, it will not appear for configuration. Uneditable options indicate mandatory items that cannot be deselected.

Figure 34: Object properties for a playbook

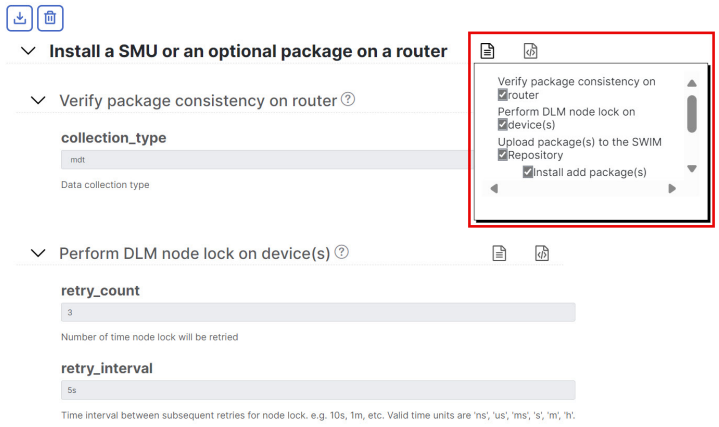

Figure 35: Object properties for a play

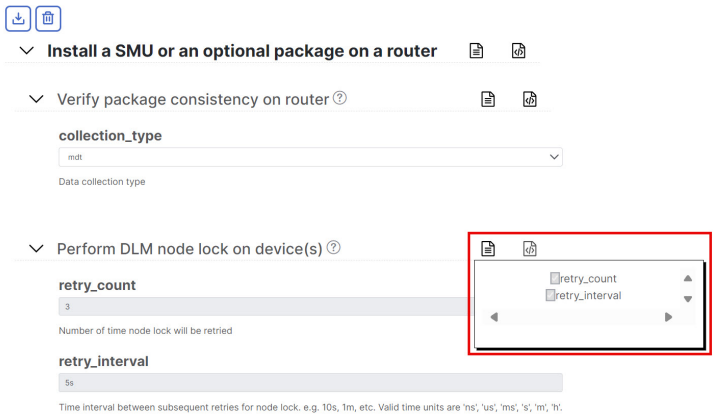

- Click ⊕ to add additional instances of a particular parameter, if required for the playbook you are running. Click 🗑 to delete instances added in this way.
- Click 🗑 to clear all the parameter values entered so far.

Step 6 With the parameter values set, click **Next**. The **Execution policy** window appears.

Step 7 Choose **Single stepping**. The **Execution policy** window displays additional features to customize the job:

Figure 36: Execution policy

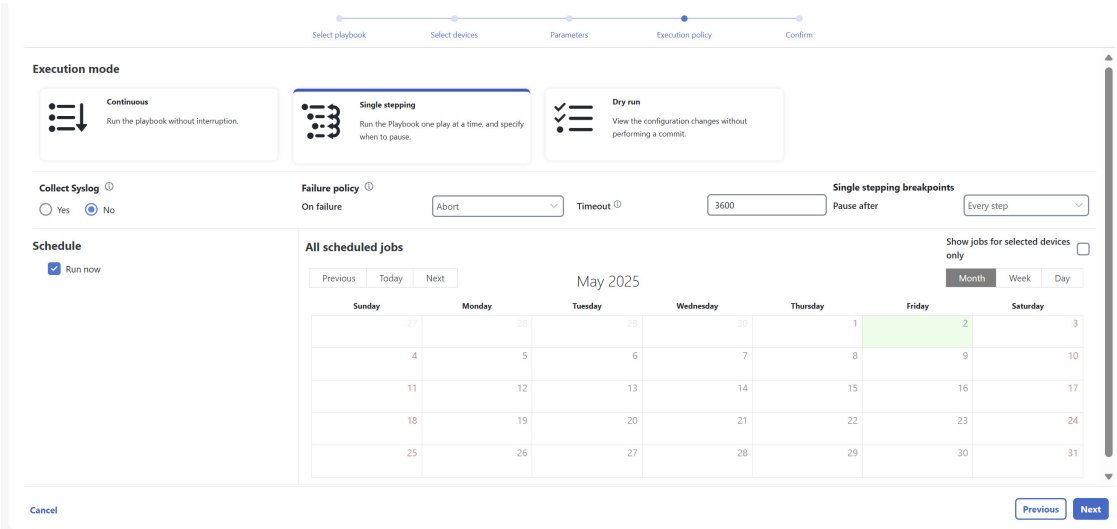

- Under **Collect syslog**, click **Yes** if you want syslogs to be collected during and immediately after the run and **No** if you do not. **Yes** is the default selection only if you have a syslog provider configured.
- From the **Failure policy** dropdown, select:
  - **Abort** to abort the entire run without rolling back any changes, if a failure occurs at any point. This is the default. Any configuration changes made up to the point of failure will not be rolled back.

- **Pause** to pause the run and allow you to decide how to handle the failure. This pause will be in addition to any breakpoints you set using the **Single stepping breakpoints** dropdown.
- **Complete roll back** to abort the entire run and roll back all configuration changes made.

- In the **Schedule** area, uncheck the default **Run now** selection to schedule the job for a later time. See Schedule playbook runs, on page 56 for help on using the **Schedule** area features.

Step 8 From the **Single stepping breakpoints** dropdown, select either

- **Every step** to pause automatically after every step in the playbook.
- **Customize** to select the steps where you want the playbook to pause.

If you select **Customize**, the **Customize check point** pop-up displays a list of all the plays in the playbook, with a ⏸ at the step between each play. Click the ⏸ at each step where you want to set a breakpoint. When you are finished, click **Done**.

Step 9 Click **Next**. The **Review your job** window appears, displaying a summary of your choices: playbook, devices, parameters, and execution policy. In this window:

- You must provide a relevant **Name** for the job.
- You can enter labels for your job using the **Labels** field.
- You can click on any **Change** links in the **Review your job** window summary to modify your choices.

Step 10 (Optional) Enter the device credentials (name and password).

Note
This step is applicable only if **Credential prompt** is enabled in the Change Automation settings. For more information, see Verify installation and configure system settings, on page 5.

Step 11 When you are ready to continue, click **Run playbook**.

Step 12 At the confirmation prompt, click **Confirm**. Click **View job set** to view the status of the current job. The job details include job status, job set tags, the title of the selected playbook, execution parameters and policy, last updated date, and updated comments (if any).

Step 13 While the run is executing, the **Running** tile at the top of the window will change to **Paused** for each step at which you have set a breakpoint. Your choices at each pause will be displayed as buttons below the tiles:

- Click **Resume** to resume running from this point, with no changes. The **Resume** request includes the runtime parameters from the previous step; you can edit these, as needed, later.
- Click **Roll back** to roll back any changes made so far. You can choose how far to rollback:
  - Click **Complete roll back** to roll back all changes to the start of the playbook run. Once you have rolled back to the start, you can choose to **Resume** from that point, **Abort** the run entirely, or **Edit runtime parameters** of the run.
  - Click **Select roll back point** to roll back changes to your selected step. All the previous steps will have a roll back point icon next to them. Click this icon for the step to which you want to roll back. Once you have selected the step, you can choose to **Resume** from that step, **Roll back** further, **Abort** the run entirely or **Edit runtime parameters**.
- Click **Abort** to abort the run entirely. No changes made will be rolled back.

- Click **Edit runtime parameters** to edit the parameters the run is using. You edit using a pop-up version of the **Parameters** window, just as you did in step 6. When resuming, the parameters exposed for editing are specific to the task being resumed, meaning they are not the same global parameters you defined in step 6. Most of the time, they are a subset of the global parameters. When you are finished, click **Apply**. You can then choose to **Resume** execution with the changed parameters.

Step 14 While the run is executing, you can also use the following features of the progress window:

- View the execution status of each play in the playbook in the **Maintenance** play list at the left side of the window. plays that fail are indicated with a red icon; plays that succeed are indicated with a green icon.
- See reminders of your choices in the **Playbook** and **Devices** tiles at the top of the window.
- See the current status of the run in the **Running** tile at the top of the window.
- Click **View** in the **Parameters** tile to view the run's parameters. While viewing the parameters, you can click **Download parameters** to save them in a JSON file. You will be prompted to name and save the file appropriately for your browser and operating system.
- Use the network topology in the map at the right side of the window to view the device and its connections to the rest of your network.

Step 15 After the run is complete:

- Click the **Events** tab to see success and failure messages for each step of the playbook. This can help you diagnose and correct problems with individual plays and the run as a whole.
- Click the **Syslog** tab to access syslog messages collected during and immediately after the run. If syslog collection is enabled, the tab will provide a pointer to the path on the syslog storage provider where collected syslogs are stored. If you chose not to collect syslogs, or no syslog storage provider has been configured, this tab will display a message indicating that syslog collection is disabled.
- Click the **Console** tab to see relevant commands and responses from the device consoles that took place during the run. These messages can also help with diagnostics.
- An event is created in the audit log (**Administration** > **Audit log**). The audit log includes details like the name of the playbook, the user who ran the playbook, and the commit label, if present.

Run playbooks in continuous mode

Continuous execution mode is the standard way to run playbooks. Configuration changes are committed to the device during the run, with no checks or delays except those programmed for system resets or other purposes. The run continues until it succeeds or fails. If it fails, you can use the run's Failure Policy to abort, rollback all changes made to the device, or pause execution at the failure point.

It is always good practice to perform a dry run and verify the configuration changes before committing to a continuous run (see Perform a dry run of a playbook, on page 38). You can also run the playbook in single-stepping mode, which will allow you to pause execution after any play you select, abort and rollback changes as needed, and even change runtime parameters in the middle of the run (see Run playbooks in single stepping mode, on page 44).

Procedure

Step 1 From the main menu, choose **Network Automation** > **Run playbook**.

Step 2 In the **Available playbooks** list on the left, click on the playbook you want to run. On the right, the window displays the playbook name, hardware and software compatibility information, and descriptions for all the plays in the selected playbook.

Figure 37: Select playbook

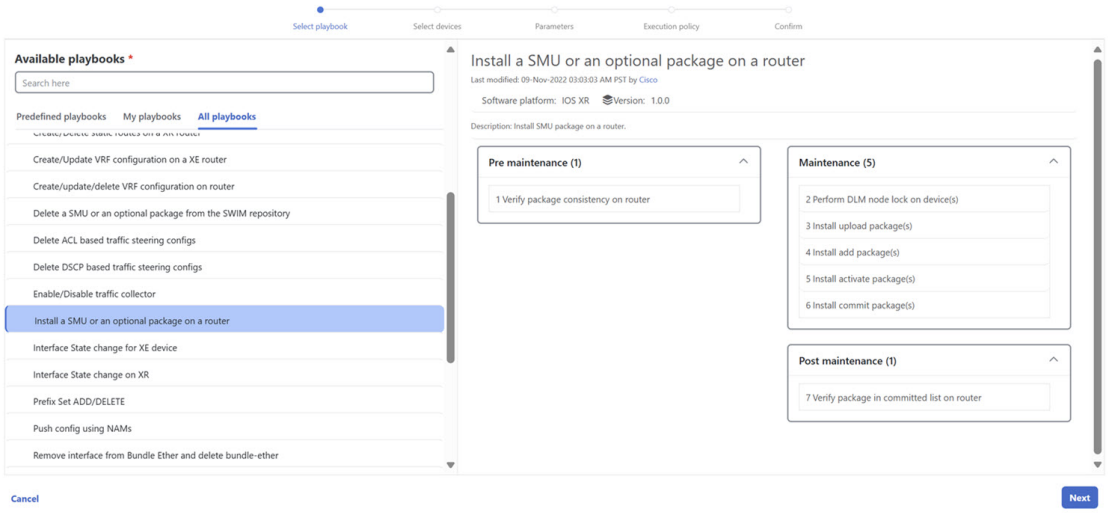

Step 3 Click **Next**. The **Select devices** window appears.

Figure 38: Select devices

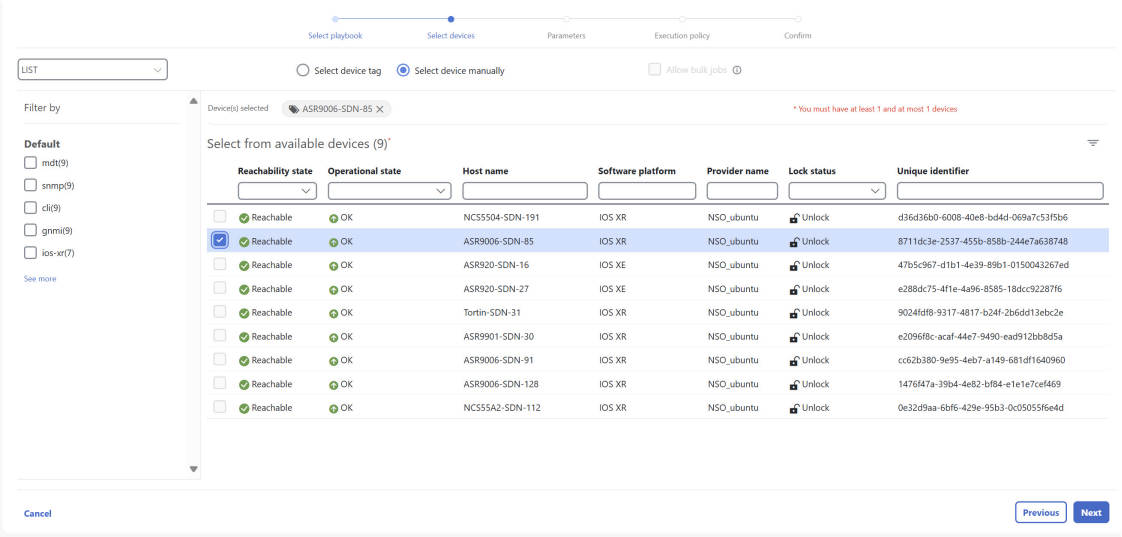

Using this window:

- You can toggle between the table view and topology map view by clicking and selecting the relevant option in the drop-down button on the upper left corner of the window. Choose **Select devices from list** or **Select devices from map** to select the table view or topology map view, respectively. By default, the table view is displayed.
- With the topology map view displayed, you can toggle between the map's geographical and logical views by clicking on the  or the  . You can also zoom, display bandwidth utilization, and change logical view layouts as you do with the topology map you see when you select **Topology** from the main menu.
- You can select the devices manually or using tags. The **Select device tag** option targets you to select the relevant tag instead of devices from the table and all devices associated with the relevant tag are selected. The **Select device manually** option allows you to select the devices from the list using quick and advanced filters and filter by tags on the left. Hover the mouse pointer over the  icon next to the options for more information. You can also view the selection criteria such as number of devices required for the selected playbook.

  Note
  If you are a non-admin user and selecting the devices manually, make a note of the following:

  - The devices on which you want to run the playbook must belong to a Device Access Group and you must have access to this Device Access Group. For more information on Device Access Groups and associating a user with a Device Access Group, see the Manage Device Access Groups section in the *Cisco Crosswork Network Controller Administration* guide.
  - If your role is associated with an empty Device Access Group, you will receive an error message.
  - If your role is associated with multiple Device Access Groups and if the device belongs to any one of these Device Access Groups, then you can run the playbook on this device. If the device does not belong to any of your Device Access Groups, the operation fails.
  - If you are selecting multiple devices (using **Allow bulk jobs** option or using tags) and if any devices does not have access, an error message appears stating that this list of devices does not have access to run the playbook.

- In the **Select device manually** selection mode, you can check the **Allow bulk jobs** check box to select multiple devices and run the selected playbook simultaneously. Based on your selection, the system creates a static group of multiple jobs. Hover the mouse pointer over the  icon next to the check box for more information. There is no limit to the number of devices you can select for a bulk job.

  Note
  Allow Bulk Jobs option is enabled for playbooks that can be executed on a single device.

Step 4 Click **Next**. The **Parameters** window appears.

Step 5 In the fields provided in the **Parameters** window, enter the playbook parameter values to use for this run.

Figure 39: Parameters

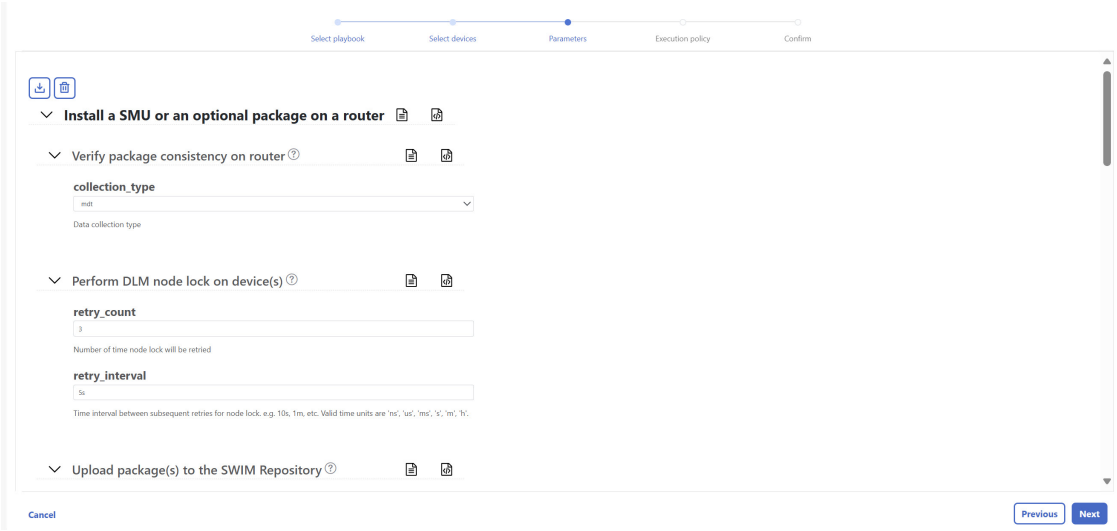

With the **Parameters** window displayed, you can also:

- Click to upload a JSON file with the parameter values you want. You will be prompted to navigate to the JSON parameters file you have previously prepared (or downloaded from a previous playbook run) and then upload it as appropriate for your browser and operating system.
- Click to enter the parameter values in JSON format. A popup text window displays the full list of JSON parameters with empty values in quotes. Edit the values, and when you are finished, click **Save**.

  Figure 40: Edit JSON

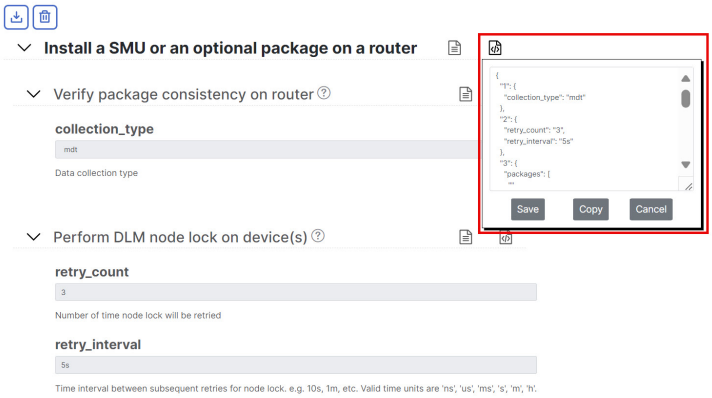

- Click to select the items you want to configure. A pop-up text window will appear, displaying the complete list of plays or parameters that can be configured. For the playbook, it will show a list of plays and, for each selected play, it will display a list of parameters that can be configured for that play. If you unselect any item, it will not appear for configuration. Uneditable options indicate mandatory items that cannot be deselected.

Figure 41: Object properties for a playbook

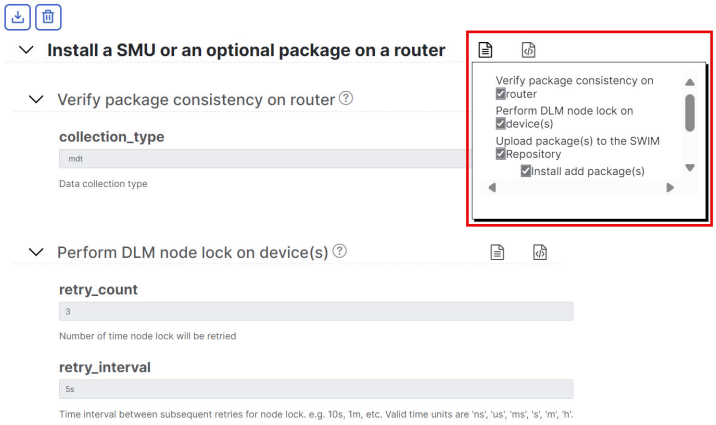

Figure 42: Object properties for a play

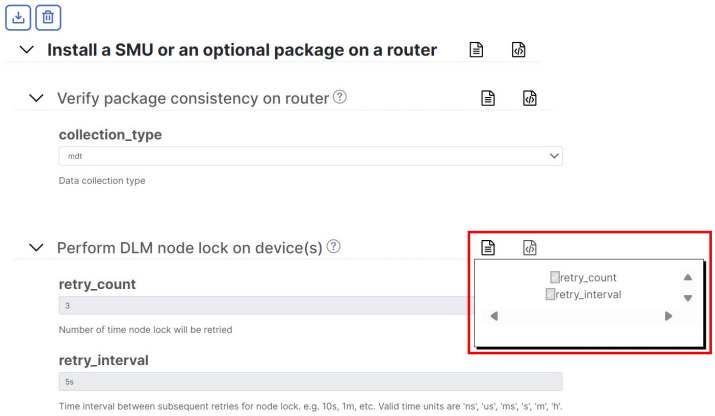

- Click ⊕ to add additional instances of a particular parameter, if required for the playbook you are running. Click 🗑 to delete instances added in this way.
- Click 🗑 to clear all the parameter values entered so far.

Step 6 With the parameter values set, click **Next**. The **Execution policy** window appears.

Step 7 Choose **Continuous**. The **Execution policy** window displays additional features to customize the job:

Figure 43: Execution policy

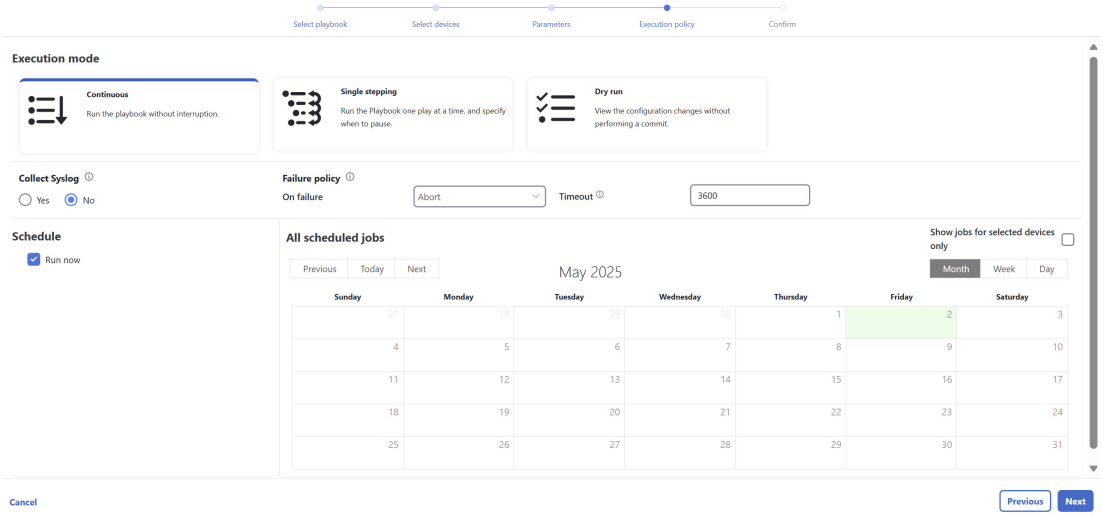

- Under **Collect syslog**, click **Yes** if you want syslogs to be collected during and immediately after the run, and **No** if you do not. **Yes** is the default selection only if you have a syslog provider configured.
- From the **Failure policy** dropdown, select:
  - **Abort** to abort the entire run without rolling back any changes, if a failure occurs at any point. This is the default. Any configuration changes made up to the point of failure will not be rolled back.
  - **Pause** to pause the run and allow you to decide how to handle the failure.
  - **Complete roll back** to abort the entire run and roll back all configuration changes made.
- In the **Schedule** area, uncheck the default **Run now** selection to schedule the job for a later time. See Schedule playbook runs, on page 56 for help on using the **Schedule** area features.

Step 8 Click **Next**. The **Review your job** window appears, displaying a summary of your choices: playbook, devices, parameters, and execution policy. In this window:

- You must provide a relevant **Name** for the job.
- You can enter labels for your job using the **Labels** field.
- You can click on the **Change** links in the **Review your job** window summary to modify your choices.

Step 9 (Optional) Enter the device credentials (name and password).

Note
This step is applicable only if **Credential prompt** is enabled in the Change Automation settings. For more information, see Verify installation and configure system settings, on page 5.

Step 10 When you are ready to continue, click **Run playbook**.

Step 11 At the confirmation prompt, click **Confirm**. Click **View job set** to view the status of the current job. The job details include information such as job status, job set tags, the title of the selected playbook, execution parameters, and policy, last updated date, and update comments (if any).

Step 12 While the run is executing, the **Running** tile at the top of the window will change to **Paused** if you chose a **Failure policy** of **Pause**. Your choices will be displayed as buttons below the tiles:

- Click **Resume** to resume running from this point, with no changes.
- Click **Roll back** to roll back any changes made so far.
- Click **Abort** to abort the run entirely. No changes made will be rolled back.

Step 13 While the run is executing, you can also use the following features of the progress window:

- View the execution status of each play in the playbook in the **Maintenance** play list at the left side of the window. Plays that fail are indicated with a red icon; plays that succeed are indicated with a green icon.
- See reminders of your choices in the **Playbook** and **Devices** tiles at the top of the window.
- See the current status of the run in the **Running** tile at the top of the window.
- Click **View** in the **Parameters** tile to view the run's parameters. While viewing the parameters, you can click **Download parameters** to save them in a JSON file. You will be prompted to name and save the file as appropriate for your browser and operating system.
- Use the network topology in the map at the right side of the window to view the device and its connections to the rest of your network.

Step 14 After the run is complete:

- Click the **Events** tab to see success and failure messages for each step of the playbook. This can help you diagnose and correct problems with individual plays and the run as a whole.
- Click the **Syslog** tab to access syslog messages collected during and immediately after the run. If syslog collection is enabled, the tab will provide a pointer to the path on the syslog storage provider where collected syslogs are stored. If you chose not to collect syslogs, or no syslog storage provider has been configured, this tab will display a message indicating that syslog collection is disabled.
- Click the **Console** tab to see relevant commands and responses from the device consoles that took place during the run. These messages can also help with diagnostics.
- An event is created in the audit log (**Administration** > **Audit log**). The audit log includes details like the name of the playbook, the user who ran the playbook, and the commit label, if present.

Schedule playbook runs

The Change Automation application's **Execution mode** window allows you to schedule future playbook runs as jobs and view all the jobs that have been scheduled. Use the **Schedule** area on the left to schedule a job. Use the **All scheduled jobs** area on the right to view scheduled jobs on the calendar.

Note **Playbook job scheduling** is only available if enabled when Change Automation was installed and initially configured. For more information, see Verify installation and configure system settings, on page 5. To change this setting, you must uninstall and then reinstall Change Automation.

Note If you are a non-admin user, ensure you have access to the Schedule playbook task. You cannot schedule playbooks without this task.

Prerequisites:

Ensure that **Playbook job scheduling** is enabled on the Device Override Credentials page. For more information, see Verify installation and configure system settings, on page 5.

To enable the task permission, do the following:

1. Go to **Administration** > **Users and Roles** > **Roles**.
2. Under the **Roles** pane, select the role for which you want to grant the access.
3. Under the **Task permissions** tab, enable the **Schedule playbook** check box and click **Save**.

The **Execution mode** window's scheduling features are only displayed when you have chosen to run a playbook in continuous or single-stepping mode. You cannot schedule a dry run of a playbook.

Figure 44: Execution mode scheduling features

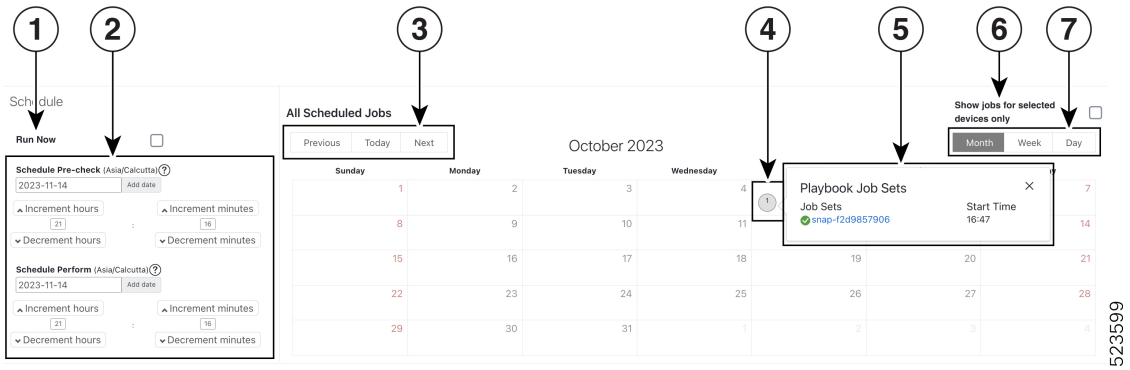

Item	Description
1	**Run now**: Running playbooks immediately is the default for continuous and single-stepping execution modes. To schedule a run for a future time and date, you must uncheck this box.
2	**Schedule selectors**: Use these fields to select the future time and date when the playbook runs. Although it is the default for the Pre-Maintenance and Maintenance phases of a scheduled playbook to start simultaneously, you can use the upper **Schedule pre-check** and lower **Schedule perform** fields to schedule the start of Pre-Maintenance and the start of Maintenance independently. The **Schedule perform** time must always be greater than or equal to the **Schedule pre-check** time.
3	**Previous/Today/Next selectors**: Use these three selectors with the **Month/Week/Day** selectors to focus the calendar's display of scheduled jobs on the time range in which you are interested. For example: To show only those jobs scheduled for next week, click **Next** and **Week**.
4	**Job icons**: Red, numbered icons in the squares representing each calendar date show how many jobs are scheduled for that date. Yellow circle icons represent each scheduled job.

Item	Description
5	**Job details pop-up**: Hover your mouse cursor over a yellow circle icon to see the details for the scheduled job represented by that icon. The pop-up shows the execution ID of the job and the name of the playbook to be run.
6	**Show jobs for selected devices only**: Check this box to restrict the calendar display to only jobs scheduled to run on the devices you have already selected. This is a handy way to see if the schedule you plan for your playbook run conflicts with other scheduled jobs on the same devices.
7	**Month/Week/Day selectors**: Use these three selectors with the **Previous/Today/Next** selectors to focus the calendar's display of scheduled jobs on the time range in which you are interested. For example: To show only those jobs scheduled for last month, click **Last** and **Month**.

Note Change Automation playbooks have a **`mop_timeout`** parameter, which is a user-specified input needed to schedule any playbook.

If you are scheduling a playbook with **Failure policy** set to **`Complete Roll Back`**, first dry run the play and note the time taken. Then, add a buffer time (for example, 10 minutes) to the time taken during the dry run. After that, double the time value and enter it to the **`mop_timeout`** parameter, as it can take as much time to roll back the playbook as it takes to run it until the last step. Without sufficient **`mop_timeout`**, the playbook can end up incomplete (in between transitions) if the timeout gets triggered while rollback is in progress. If this happens, you have to revert the changes manually or create a playbook with the changes you want to revert.

View or abort playbook jobs

The **Automation Job History** window lets you click on any individual job in the list to see that job's detailed execution progress panel. This panel displays the name of the playbook, its plays, the devices it ran on, the parameters used, and all events, Syslog, console, and other messages. These details are helpful when diagnosing failures.

The **Automation Job History** window also allows you to abort *running* jobs.

You can also navigate to **Automation Job History** window from the **Jobs** panel in the Change Automation **Dashboard**.

Before you begin

A user must have the permission for specific playbook label to run or abort a playbook. For more information on assigning playbooks to specific roles, see Assign playbooks to specific roles, on page 36.

Procedure

Step 1 From the main menu, select **Network Automation** > **Automation Job History**. The **Automation Job History** window displays a list of Job Sets.

Figure 45: Automation Job History

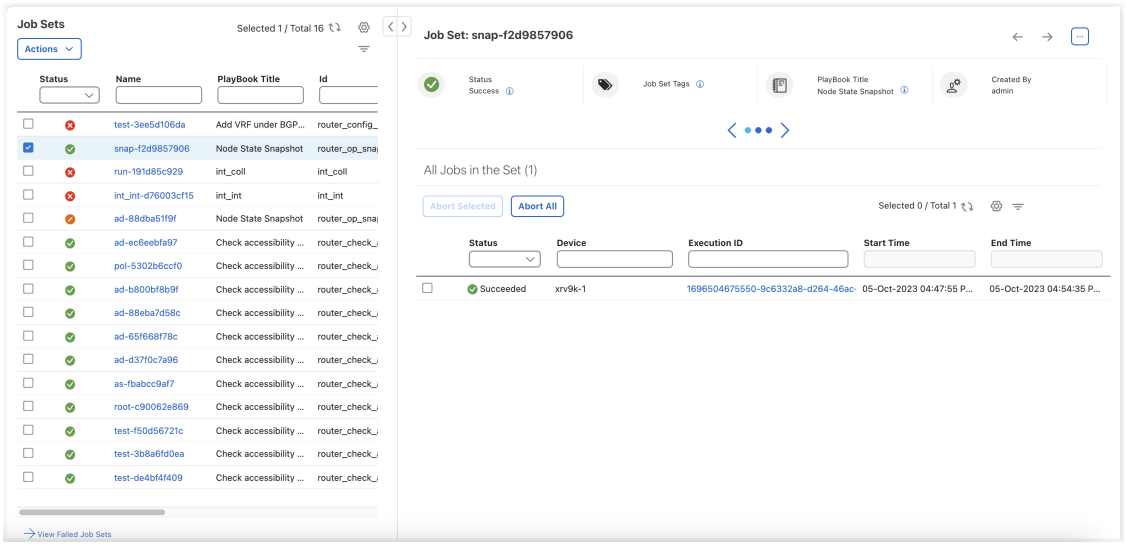

The list in **Automation Job History** window is sorted by the last update time, with running or most recently executed jobs at the top. You can apply quick or advanced filters to the table as you would with columns in other table windows.

Step 2 To view information about a playbook job, click the relevant job ID checkbox on the left. The job's status and execution details are displayed on the right side. Click on the ⓘ icon next to each detail to get more information about the selected job set.

Step 3 You can abort a job set in running, paused or scheduled status, as follows:

- To abort a specific job, click the check box next to it and then click **Abort selected**.
- To abort all jobs immediately, click **Abort all**.

When prompted, click **Confirm**. Jobs currently running, paused, or scheduled will abort once the current task has been completed.

Use the change automation dashboard

The Change Automation **Dashboard** window (shown in the figure below) lets you view all playbook-related activity and initiate playbook runs. It displays the total number of playbooks, the playbook Jobs Calendar, the most recently run playbook jobs, and the same network topology map you see when you select **Topology** from the main menu.

To view the Change Automation **Dashboard** window, select **Network Automation** > **Dashboard**.

Figure 46: Change Automation Dashboard Window

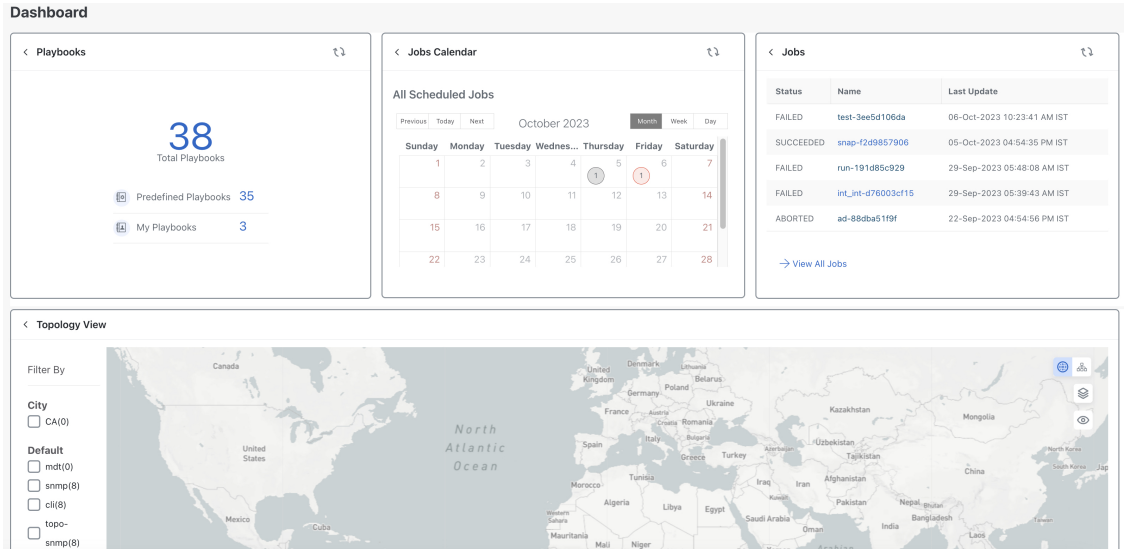

The **Playbooks** tile displays the total number of playbooks (pre-defined and custom). Clicking on a specific number displays all the playbooks that correspond to the selected category:

- **Total playbooks** indicate the total number of pre-defined and user-created playbooks (My playbooks) in the system.
- **Predefined playbooks** indicate the number of pre-defined playbooks that exist in the system.
- **My playbooks** indicate the number of custom playbooks created by the current user.

Creating playbooks does not use a license. The license count is incremented only upon the first execution of a playbook (pre-defined or user-created), irrespective of whether the playbook runs successfully. Subsequent execution of the playbook does not increment the license count.

The **Jobs calendar** tile displays a calendar (month, week, day) with the number of job sets executed on a given day marked in a circle against the corresponding date. Clicking on the number displays a dialog box with the names of the playbook job sets and their execution time. Click the desired job set to view the execution details.

The color of the circle indicates the overall status of the job sets:

- A **red** circle indicates at least one job set with **Failed** status among the day's overall job sets.
- A **gray** circle indicates that all job sets are in **Scheduled** or **Running** status.
- A **blue** circle indicates at least one critical job set in **Recovered** status among the day's overall job sets.
- A **green** circle indicates most of the playbooks are in success state. Clicking on it displays all the jobs that are **Recovered**, **Scheduled**, or **Running**.

The **View All Jobs** link on the **Jobs** tile gives you direct access to the Change Automation **Automation Job History**.

Troubleshoot change automation

The following table describes issues you may encounter when using the Change Automation application and their solutions or workarounds.

Table 1: Change Automation Troubleshooting

Issue	Solution
playbook run fails with messages indicating that Cisco Network Services Orchestrator (Cisco NSO) and the target device are out of sync or communication. Message text varies but may include "device out of sync," "NC client timeout," and other text indicating connectivity or sync issues between Cisco NSO and the device.	Ensure that the playbook does not include a sync operation. Get the device and Cisco NSO back in sync, and then re-run the playbook. Alternatively, you can create a new playbook that includes a sync operation to avoid future problems.
playbook run fails with "access error" messages indicating failure "to set device override credentials in NSO".	Ensure that "admin" is a member of the **ncsadmin** user group in Cisco NSO.

"Failed to end NSO transaction, 500:fatal:YClientError: Failed to send RPC:" error is displayed while running the playbook.	Include the below settings in the Cisco NSO configuration file (`ncs.conf`):
	`<ssh>`
	`<client-alive-interval>infinity</client-alive-interval>`
	`<client-alive-count-max>5</client-alive-count-max>`
	`</ssh>`
	Note
	This configuration could increase the load on Cisco NSO, so it is better to do it only when necessary.

Playbook aborted due to failure in locking the device nodes.	In the Devices window, select the relevant devices and clear the lock by moving the device to DOWN and then UP. Go to **Administration** > **Crosswork Manager**, click the **Change Automation** tile, and restart the robot-nca process. Once the protocols are reachable, you can schedule to run a new playbook.
SMU install fails at "Verify package in committed list on router".	Instead of using the tar.gz file in the packages field under Verify package in the committed list on router sub-option, use the committed package name to verify the package.

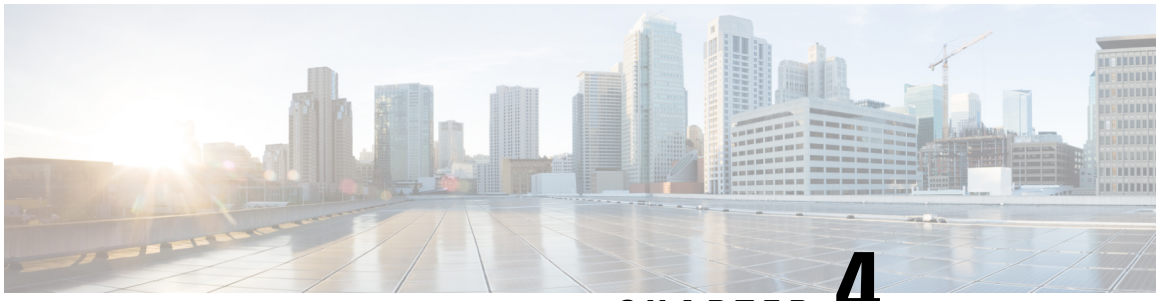

CHAPTER 4

Monitor Network Health and KPIs

This section contains the following topics:

Health Insights overview

Health Insights is a network health application that:

- performs real-time key performance indicator (KPI) monitoring, analytics, and alerting and aids in troubleshooting.
- builds dynamic detection and analytics modules that allow operators to monitor and alert network events with user-defined logic.
- provides prebuilt KPIs that are based on Model-Driven Telemetry (MDT), SNMP-based telemetry, or GNMI/Openconfig based telemetry collection.

The Health Insights Recommendation Engine uses data mining to analyze your network and recommends which telemetry paths you should enable and monitor.

Note For MDT-based KPIs, Crossword Network Controller pushes the KPI configuration down to the device. For SNMP, CLI, and GNMI-based KPIs, the operator must have the device configured to respond to a request for telemetry data.

Important Due to the additional data collection tasks required, Health Insights requires the use of extended Data Gateways.

The following high-level example gives a basic view of how Health Insights interacts with the other Crosswork Network Controller components:

1. Health Insights detects an anomaly: The optical bit error rate that you are monitoring on each of the links in your network suddenly increases.
2. Change Automation playbooks automate remediation: Switch to the backup link immediately. Restore service. Open a ticket (manually initiated by the user). Alert the network engineer.

Health Insights is configured to gather the link bandwidth usage data for device links. After a time period, it establishes a performance baseline for each link. If a link deviates from its baseline causing an alert to be generated, Health Insights detects it and you can then go and run the Playbook to reconfigure the network to resolve the issue.

The complexity of the interaction will depend on the type of anomaly, how it is detected, and the playbooks you choose to use to remediate it. You can orchestrate any form of network remediation using Change Automation playbooks, helping you to close the loop on problem resolution and maximize network performance.

Health Insights collects telemetry data from devices and stores it for the last 72 hours in a time-series database. This data is used for real-time KPI monitoring, analytics, and generating alerts. Triggered alerts are stored and retained for 30 days in the same database, and the messages showing the alerts' duration are displayed at the top of the Device/KPI view in the Alert dashboard.

List of Health Insights KPIs

This section lists the prebuilt Health Insights KPIs supplied with Health Insights application.

Supported protocols

The target device(s) must support the form of telemetry used by the KPI either SNMP, GNMI, or MDT. The application validates for a match between KPI and device telemetry capabilities.

Definition of the protocols:

- Model-Driven Telemetry (MDT): Model-driven telemetry provides a mechanism to stream operational data from device as defined in the YANG model(s) to a data collector.
- gRPC Network Management Interface (gNMI): gNMI provides the mechanism to install, manipulate, and delete the configuration of network devices, and also to view operational data.
- Simple Network Management Protocol (SNMP): SNMP is an IP protocol for collecting and organizing information about managed devices on IP networks and for modifying that information to change device behavior.
- Command Line Interface (CLI): CLI is used in network device management.

Health Insights uses either MDT or gNMI protocols but the device supports both. gNMI is a preferred default.

Note In Crosswork Network Controller version 7.1, new sensor paths have been introduced for the Layer3-Routing, QoS, and Layer2-Traffic (Openconfig-interfaces) KPI categories to ensure compatibility with Cisco IOS XR devices running version 24.1.1 and above.

- For devices running Cisco IOS XR version 7.9.21 and earlier, the Layer3-Routing-deprecated sensor paths will remain temporarily supported.
- For devices running Cisco IOS XR version 24.1.1 and above, use the Layer3-Routing KPIs with the new sensor paths designed for seamless integration with newer devices.

Users must select the appropriate sensor paths based on the IOS XR version of their device.

Note When upgrading from an older version of Crosswork Network Controller to 7.1, any KPI profiles containing these KPIs will be disabled during migration. To apply the new sensor paths, users must manually re-enable the KPI profiles after the upgrade.

Table 2: Health Insights KPIs

KPI Name	Description	Alerting	Protocol
Basics			
Device uptime	Monitors device uptime.	Low Single Threshold	MDT, gNMI
CPU			
CPU threshold	Monitors CPU usage across route policies and line cards on routers. Generates an alert when CPU utilization exceeds the configured threshold	Two-Level Threshold	MDT, gNMI
CPU utilization	Monitors CPU usage across route policies and line cards on routers. Generates an alert when CPU utilization is unusual.	Standard Deviation	MDT, gNMI
Dataplane-Counters			
CEF drops	Monitors CEF drop counters and baseline. Generates an alert for an unusual number of drops.	Rate Change	MDT, gNMI
File System			
Filesystem Utilization	Monitors filesystem usage on active route processor and generates an alert when filesystem utilization exceeds the configured threshold.	Two-Level Threshold	CLI
IPSLA			
IP SLA UDP echo RTT	Monitors IP SLA UDP echo RTT. Generates an alert when unusual RTT values occur.	Standard Deviation	MDT, gNMI

KPI Name	Description	Alerting	Protocol
IP SLA UDP jitter monitoring	Monitors IP SLA UDP jitter. Generates an alert when an abnormal UDP jitter occurs.	Standard Deviation	MDT, gNMI
LLDP			
LLDP neighbors	Monitors LLDP neighbors.	No Alert	MDT, gNMI
Layer 1-Optics			
Layer 1 optical alarms	Monitors per-port optical alarms (current and past).	No Alert	MDT, gNMI
Layer 1 optical errors	Monitors per-port Layer 1 errors. Generates an alert when error rates exceed the configured threshold.	Rate Change	MDT, gNMI
Layer 1 optical FEC errors	Monitors per-port optical FEC errors. Generates an alert when FEC errors exceed the configured threshold.	Rate Change	MDT, gNMI
Layer 1 optical power	Monitors per-port optical power.	No Alert	MDT, gNMI
Layer 1 optical temperature	Monitors per-port optical temperature.	No Alert	MDT, gNMI
Layer 1 optical voltage	Monitors per-port optical voltage.	No Alert	MDT, gNMI
Layer 1-Traffic			
Ethernet port error counters	Monitors port transmit and receive error counters.	Rate Change	MDT, gNMI
Ethernet port packet size distribution	Monitors port transmit and receive packet size distributions.	No Alert	MDT, gNMI
Ethernet port packet statistics	Monitors port transmit and receive packet statistics.	Standard Deviation of Rate Change	MDT, gNMI
Layer2-Interface			
Interface flap detection	Monitors interface flaps and alerts when flap count reaches set threshold.	Two-Level Threshold	MDT, gNMI
Line state	Monitors interface line states.	No Alert	MDT, gNMI
Layer 2-Traffic			
Interface bandwidth monitor	Monitors bandwidth utilization across all interfaces on a router. Generates an alert when bandwidth exceeds the configured threshold.	Two-Level Threshold	MDT, gNMI
Interface packet error counters	Monitors interface transmit and receive error counters. Generates an alert when unusual error rates occur.	Rate Change	MDT, gNMI

KPI Name	Description	Alerting	Protocol
Interface packet error counters (Openconfig)	Monitors interface error counters; generates an alert when unusual error rates occur. This KPI uses openconfig-interfaces YANG model. **Note** In Crosswork Network Controller version 7.1, the sensor paths for this KPI have changed. When upgrading from an older version, KPI profiles that include this KPI will be disabled and must be manually re-enabled to apply the new sensor paths.	Rate Change	gNMI
Interface packet counters	Monitors interface transmit and receive counters. Generates an alert when unusual traffic rates occur.	No Alert	MDT, gNMI
Interface rate counters	Monitors interface statistics as rate counters. Generates an alert when unusual traffic rates occur.	Standard Deviation	MDT, gNMI
Interface rate counters (Openconfig)	Monitors interface statistics (such as rate counters), and generates an alert when unusual traffic rates occur. **Note** In Crosswork Network Controller version 7.1, the sensor paths for this KPI have changed. When upgrading from an older version, KPI profiles that include this KPI will be disabled and must be manually re-enabled to apply the new sensor paths.	Rate Change	gNMI
SNMP interface packet error counters	Monitors interface transmit and receive error counters. Generates an alert when unusual error rates occur.	No Alert	SNMP
SNMP interface packet counters	Monitors interface transmit and receive counters.	No Alert	SNMP
SNMP interface rate counters	Monitors interface statistics as rate counters. Generates an alert when unusual traffic rates occur.	Standard Deviation Rate of Change	SNMP
Layer 3-Routing **Note** In Crosswork Network Controller version 7.1, sensor paths for Layer 3-Routing KPIs have changed to support Cisco IOS XR devices running version 24.1.1 and above. When upgrading from an older version, KPI profiles that include these KPIs will be disabled and must be manually re-enabled to apply the new sensor paths.			
IPv6 RIB BGP route count	Monitors IPv6 RIB for route count and memory used by BGP. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
IPv6 RIB IS-IS route count	Monitors IPv6 RIB for route count and memory used by IS-IS. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI

KPI Name	Description	Alerting	Protocol
IPv6 RIB OSPF route count	Monitors IPv6 RIB for route count and memory used by OSPF. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB BGP route count	Monitors RIB for route count and memory used by BGP. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB connected route count	Monitors RIB for route count and memory used by connected. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB IS-IS route count	Monitors RIB for route count and memory used by IS-IS. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts)	Standard Deviation	MDT, gNMI
RIB local route count	Monitors RIB for route count and memory used by local. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB OSPF route count	Monitors RIB for route count and memory used by OSPF. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB static route count	Monitors RIB for route count and memory used by static. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 connected route count	Monitors RIBv6 for route count and memory used by connected. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 local route count	Monitors RIBv6 for route count and memory used by local. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 static route count	Monitors RIBv6 for route count and memory used by static. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
Layer 3-Routing-deprecated **Note** The sensor paths for these KPIs are compatible with Cisco IOS XR devices running version 7.9.21 and earlier.			
IPv6 RIB BGP route count deprecated	Monitors IPv6 RIB for route count and memory used by BGP. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
IPv6 RIB IS-IS route count deprecated	Monitors IPv6 RIB for route count and memory used by IS-IS. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI

KPI Name	Description	Alerting	Protocol
IPv6 RIB OSPF route count deprecated	Monitors IPv6 RIB for route count and memory used by OSPF. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB BGP route count deprecated	Monitors RIB for route count and memory used by BGP. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB connected route count deprecated	Monitors RIB for route count and memory used by connected. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB IS-IS route count deprecated	Monitors RIB for route count and memory used by IS-IS. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB local route count deprecated	Monitors RIB for route count and memory used by local. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB OSPF route count deprecated	Monitors RIB for route count and memory used by OSPF. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIB static route count deprecated	Monitors RIB for route count and memory used by static. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 connected route count deprecated	Monitors RIBv6 for route count and memory used by connected. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 local route count deprecated	Monitors RIBv6 for route count and memory used by local. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
RIBv6 static route count deprecated	Monitors RIBv6 for route count and memory used by static. Generates an alert when an anomaly is detected (such as significant increase or decrease in route counts).	Standard Deviation	MDT, gNMI
Layer 3-Traffic			
Interface counters by protocol	Monitors interface statistics (such as incoming and outgoing packets or byte counters) organized by protocol.	No Alert	MDT, gNMI
Memory			
Memory utilization	Monitors memory usage across route processor and line cards on routers. Generates an alert when memory utilization is unusual.	Standard Deviation	MDT, gNMI

KPI Name	Description	Alerting	Protocol
Memory utilization (cXR)	Monitors memory usage across route processor and line cards on classic XR devices. Generates an alert when memory utilization is unusual.	Standard Deviation	MDT, gNMI
Protocol-ISIS			
ISIS neighbor summary	Monitors ISIS neighbor summaries for changes in neighbor status.	No Alert	MDT, gNMI
QoS **Note** In Crosswork Network Controller version 7.1, sensor paths for QoS KPIs have changed. When upgrading from an older version, KPI profiles that include these KPIs will be disabled and must be manually re-enabled to apply the new sensor paths.			
Interface QoS (egress)	Monitors interface QoS on the egress direction for queue statistics, queue depth, and so on.	No Alert	MDT, gNMI
Interface QoS (ingress)	Monitors interface QoS on the ingress direction for queue statistics, queue depth, and so on.	No Alert	MDT, gNMI

Manage KPIs

The Health Insights Key Performance Indicators (KPI) window gives you complete access to Cisco-supplied and user-created KPIs. You can add, edit, delete, import, and export your KPIs. You can also link your KPIs to the Change Automation application's playbooks.

To display the Health Insights Manage KPIs window, choose **Performance Alerts** > **Key Performance Indicators (KPI)** from the main menu.

Figure 47: Key Performance Indicators (KPI)

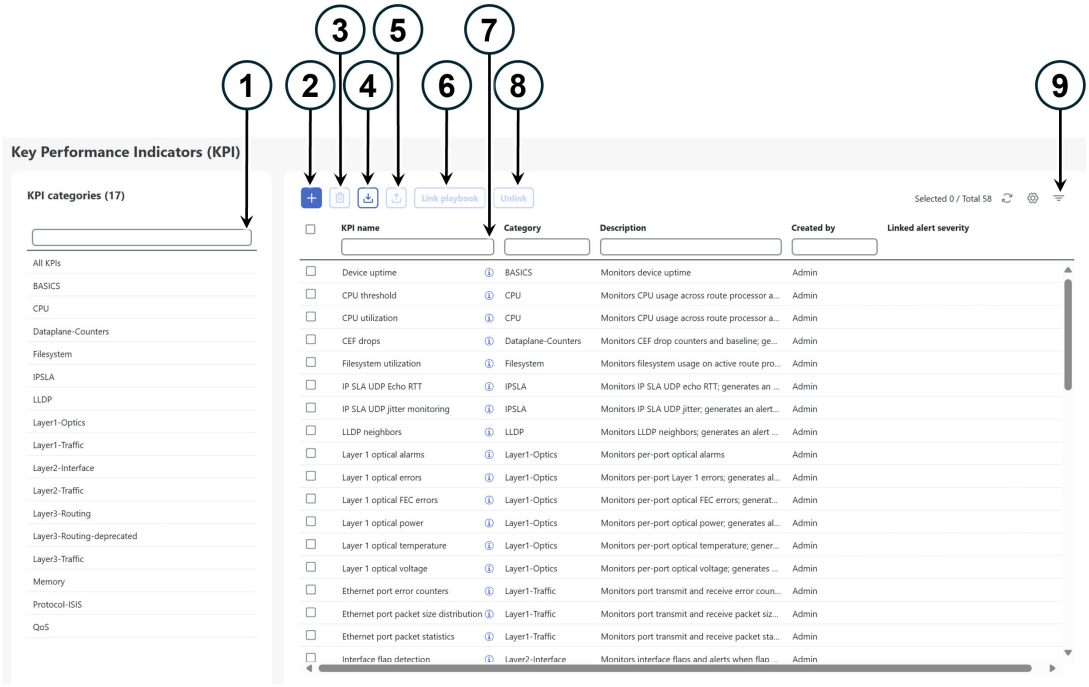

Item	Description
1	**Filter KPI Categories**: To find a KPI category, enter all or part of the KPI Category name in this field. Then click to filter the list below.
2	**Add KPIs**: Click to add a new, user-created KPI. For help with this task, see Create a new KPI, on page 73.
3	**Delete KPIs**: Select one or more existing user-created KPIs in the list and then click . You will be prompted to confirm that you want to delete the KPIs. Click **Delete** to confirm. **Note** You can delete user-created KPIs only. You cannot delete Cisco-supplied KPIs.

Item	Description
4	**Import KPIs**: Click to import new user-written or Cisco-supplied KPIs. **Note** When upgrading from an older version of CNC, it's important to consider the following: Before attempting to load a KPI, ensure that it complies with the requirements of the current release. If you try to load a KPI that was created for a previous release and is not compatible, you will receive an error message. You will be prompted to browse to the gzipped tar archive containing the KPIs to be imported. When you have selected the archive, click **OK** to begin importing it. Once imported, the new KPIs appear in the list of KPIs, with each KPI name and category assigned based on the definition in the KPI itself. In order for Health Insights to import them, KPI files must: • Be packaged as a gzip tar archive. You can include more than one KPI in a single archive; each will be imported as a separate KPI. • Have unique names and descriptions. These must not match the name or description of any Cisco-supplied KPI. If the name or description of the KPI matches an existing user-created KPI, the import will overwrite the existing KPI. • Meet other minimum requirements for Health Insights KPIs, as explained on Cisco DevNet.
5	**Export KPIs**: Select one or more existing KPIs in the list and then click to export them. Health Insights will package the exported KPIs as a single TGZ archive with a unique name. Your browser will then prompt you to save the archive to a name and location in your local file system that you select.
6	**Link Playbooks**: Select a KPI and then click Link Playbook to link it to a playbook. Linking a playbook streamlines the remediation process by importing data from the alert and using it to pre-populate the parameters the playbooks needs (such as device, interface names, and so on) to run in order when you attempt to remediate the issue. For help with this task, see Link KPIs to playbooks and run them manually, on page 75.
7	**Filter KPIs**: To find a KPI, enter all or part of the **KPI Name**, **Category**, **Description**, or **Linked Playbook** in the fields provided. The list below is automatically filtered to match your typed entry. Filtering is case-sensitive. Click to clear any filter criteria you may have set.
8	**Unlink Playbooks**: Select a KPI with a linked playbook and then click Unlink to unlink the playbook. You will be prompted to confirm that you want to unlink the playbook. Click **Unlink** to confirm.
9	**Filter**: Click to set filter criteria on one or more columns in the table.

Create a new KPI

You can create a custom KPI and enable it on the desired devices. The workflow is as follows:

1. Supply basic information, such as the KPI name and a summary description.
2. Set the KPI cadence.
3. Select a YANG module and choose sensor paths.
4. Select an alert template and set its parameters.
5. Enable the KPI on the devices.

Note Health Insights supports creating and using KPIs that use GNMI as the transport and use sensors that are based on Open Config (OC) YANG modules for collecting telemetry data (with GNMI transport). The requirements for this feature are:

- GRPC must be configured in your device.
- The device properties, while onboarding, must include GNMI under the **Capability** field, and the GNMI protocol details must be provided under the **Connectivity Details** field.
- While creating a KPI, choosing an OC YANG module supports the KPI affinity for GNMI transport, while choosing Cisco-provided YANG models provides the KPI affinity for both MDT and GNMI transports.

The GNMI transport capability is determined at runtime which is based on the following factors such as GNMI capability of the device, GNMI affinity of the KPI, and the combined capability as a set of devices in a KPI Profile.

The following steps explain how to create a KPI:

Before you begin

Make sure that the device packages for the devices you want to monitor are available in Crosswork. If they are not available, perform the Add custom packages procedure given in the *Cisco Crosswork Network Controller Administration* guide. Then continue with the steps below.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Key Performance Indicators (KPI)**. The **Key Performance Indicators (KPI)** window is displayed.

Step 2 Click the  . The **Create KPI** window opens.

Step 3 In the text fields provided, enter a unique **KPI name**, a short **KPI summary** description, and **KPI details**. The **KPI group** is preset to `User Created`.

Step 4 The **Cadence** field sets the number of seconds between data collections . Leave it at the default or use the numerical selector to choose a different value.

Step 5 In the **YANG modules** area, choose one module and one or more sensor paths from which to stream data:

a) Use the **Module** field to filter and choose the desired Cisco IOS XR YANG module.
b) Use the table fields to filter and choose the desired sensor path. When you choose a path, the leaf node gets resolved to the base encoding path. If the YANG module is hierarchical, the field names are concatenated down from the base path. Only one gather path is supported for user-created KPIs.

Note
If the devices are not listed in the default YANG modules, you can expand the device coverage. Perform the Add custom packages procedure given in the *Cisco Crosswork Network Controller Administration* guide, then continue with the subsequent steps in this procedure.

Click **Next** to display the **Select alert templates** window. You can select from the following alerting types:

- **No Alert**: The KPI gathers, tracks, and reports performance data without triggering alerts.
- **Standard Deviation**: The KPI detects spikes or drops in measured values and alerts when these values deviate some number of standard deviations away from their normal values.
- **Two-Level Threshold**: The KPI detects abnormal measured values using two custom thresholds and the ability to provide dampening intervals on the thresholds.
- **Rate Change**: The KPI detects abnormal rates of change in measured values to detect rising or falling values.

You can also use the following additional alerting types when you export and modify a prebuilt KPI to create a KPI with custom parameters:

- **Standard Deviation of Rate Change**: The KPI alerts on standard deviations of the rate of change.
- **Low Single Threshold**: The KPI alerts on a single threshold when the value falls below that threshold.
- **Direct Alarm Forwarding**: The KPI uses the alarm from the device directly, as a Health Insights KPI alert.
- **Major/Minor/Low/High Thresholds**: The KPI alerts on Major high, Minor high, Minor low, and Major low values.
- **Line State Changes**: The KPI alerts on shutdowns and flapping in line states.

Note
To build a KPI that uses data from more than one module, you can do this with KPI profiles and alert groups. For more information, see Create a new KPI profile, on page 80.

Step 6 Choose the alert template that you want to use with your new KPI: **No Alert**, **Standard Deviation**, **Two-Level Threshold** or **Rate Change**. Then click **Next** to display the **Alert parameters** window appropriate for the type of alert template you chose.

Step 7 Edit the alert template parameter values as appropriate for the template and the purpose of your KPI, as follows:

- Use the **Basic** and **Advanced** parameters dropdowns to view and edit the parameter sets you need.
- Change alert parameter numerical values using the selectors or by editing the field contents
- Change alert parameters with discrete choices using parameter field dropdowns and select each choice as needed.
- Learn more about an alert parameter: Hover your mouse cursor over the ⓘ shown next to the parameter name.
- Click the **View HI Subservice class** link to view the tick script code you are generating with your changes. The tick script code updates as you make your edits. At any time, click the **Hide HI Subservice class** to close the tick script code window.

Step 8 When you are finished making changes, click **Finish** to save the new KPI and display the **Key Performance Indicators (KPI)** window.

Link KPIs to playbooks

You can link any Health Insights KPI to one Change Automation playbook of your choice. You can run the linked playbook whenever the linked KPI raises an alert in response to the event associated with the performance indicator the KPI is monitoring. The KPI alert can be raised in response to a threshold crossing, topology changes, flapping conditions, and other parameters. These parameters vary, as appropriate, for each KPI.

Link KPIs to playbooks and run them manually

The default option for KPI-linked playbooks is for the network operator to run them manually, when an alert is displayed. Crosswork displays the linked playbooks as options, and the operator can select which playbooks to run. However, if Device Override Credentials are enabled properly, you have the option to run one or more KPI-linked playbook automatically, whenever the linked KPI raises an alert, as explained in Link KPIs to playbooks and run them automatically, on page 77.

Note You can't use this function if you haven't installed the Change Automation Crosswork application. If that's the case, Crosswork will not display the UI features that link Health Insights KPIs and Change Automation playbooks (for example, you won't see the Link Playbook icon).

You can specify the **Source** of the parameter values the linked playbooks use when you run them. When linking a playbook to a KPI alert, you can select these sources:

- **Playbook**: Use default values coded into the playbook itself
- **KPI Alert**: Use values that are taken from the alert that is raised by the linked KPI.
- **Run-time Input**: Use values that you enter only at the moment you run the playbook.

The ability to set the source of these playbook parameter values gives you flexibility in how you use the linked playbook. For example: Link the KPI **Interface flap detection**, which detects interface flapping, to the playbook **Interface state change on XR**, which can be used to set the interface up or down. Depending on circumstances, you may want to set the playbook parameters as follows:

- **Playbook**: You want to run the playbook as it normally does, so you would set the **Source** as **`Playbook`** for the *provider*, *collection_type* and *mop_timeout* parameters. In the case of the *collection_type*, you can still choose between **`telemetry`** and **`snmp`**, depending on whether you want to use MDT or SNMP to gather device data.
- **KPI alert**: You want the playbook to run only on the host device and interface affected by the flapping, which are identified in the flap-detection Alert. So set the **Source** of the playbook's *hosts* and *if_names* parameters to **`KPI Alert`**. You can then use the alert's data about the **`Producer`** device and the **`interface_name`** of the flapping interface on that device.
- **Run-time input**: You want the freedom to decide at runtime whether to bring the flapping interface up or down. So set the **Source** of the playbook parameter *admin_state* to **`Runtime Input`**. The playbook will prompt you for an **`up`** or **`down`** choice when you initiate the run.

The following figure shows what this set of choices will look like:

Figure 48: Example: Specifying parameter value sources for a linked playbook

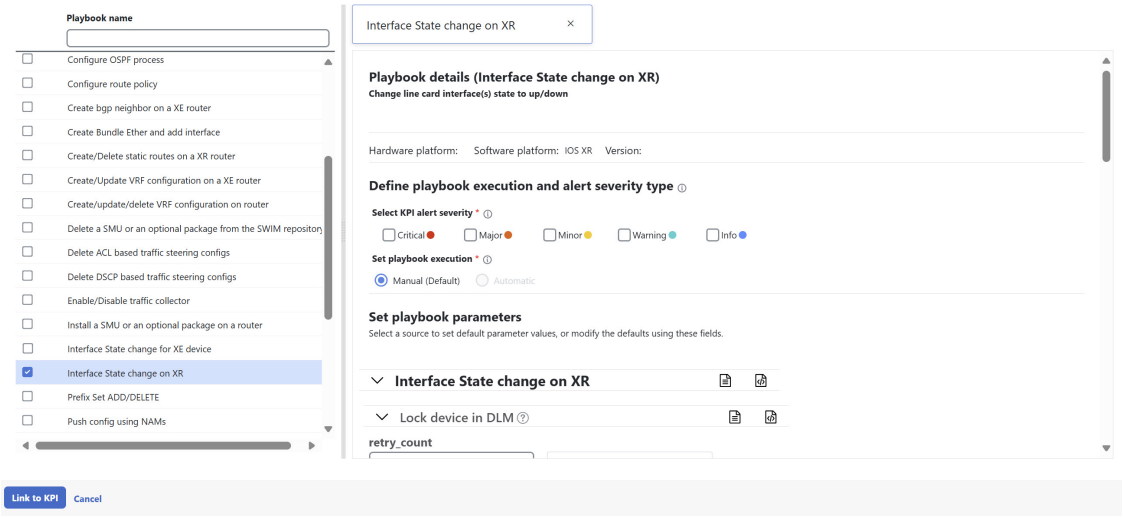

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Key Performance Indicators (KPI)**. The **Key Performance Indicators (KPI)** window opens, displaying lists of the KPI categories and the KPIs available in each category.

Step 2 Select the KPI you want to link to a playbook. You can use filters to find the KPI you want, as explained in Manage KPIs, on page 70.

Step 3 Click Link Playbook. The **Link Playbook to KPI** window opens.

Step 4 The left side of the window lists the name of the selected KPI and the playbooks appropriate for linking to it. Scroll through the list, or use the **Playbook name** field to restrict the list to just the playbooks you want to see.

Step 5 When you have found the playbook that you want to link to your chosen KPI, click the playbook name. The right side of the window will then list the **Playbook details** for the selected playbook, including:

- The **Hardware platform** and **Software platform** with which the playbook is compatible.
- The minimum software **Version** required to execute the playbook.
- The **KPI alert severity** level needed to trigger a run of this playbook. Note that, if you will be choosing multiple playbooks to run when a KPI alert is raised, be aware that playbooks do not share severity levels. If you have selected a **Critical** severity level for one playbook, you must select **Major**, **Minor**, **Warning** or **Info** severities for a second playbook, and another still for a third playbook.
- Choose the **Set playbook execution** function you want to use. **Manual** execution is the default and is recommended for most purposes. See Link KPIs to playbooks and run them automatically, on page 77 before selecting the **Automatic** execution option.
- Modify the **Set playbook parameters** default values to be used when the playbook runs. In many cases, you can select from a range of default values. You can also enter your own. These values vary widely, depending on the playbook and its purpose. For help, see the information offered on screen for the playbook you have selected.

Step 6 Verify or modify the **Source** and parameter values as needed.

Step 7 When you are finished making changes, click `Link to KPI`. The **Key Performance Indicators (KPI)** window is displayed again, this time with the linked playbooks shown next to the name of the KPI in the **Key Performance Indicators (KPIs)** list.

Step 8 If you want to run more playbooks (up to three playbooks total): Repeat steps 5 through 7 for each additional playbook you want to run when an alert for this KPI is raised.

Step 9 To change the playbook parameters linked to a given KPI, repeat steps 5 through 7 for that KPI, but this time choose the playbook whose settings you want to modify. If you have chosen multiple KPIs, you can switch among them by clicking on the playbook tiles at the top of the window. To unlink a playbook entirely, select the KPI and click Unlink.

Link KPIs to playbooks and run them automatically

In addition to running KPI-linked playbooks only at the network operator's discretion, you can choose to run one or more of your KPI-linked playbooks automatically, whenever the KPI linked to that playbook raises an alert of sufficient severity.

Note You can't use this function if you haven't installed the Change Automation application. If that's the case, Crosswork will not display the UI features that link Health Insights KPIs and Change Automation playbooks (for example, you won't see the Link Playbook icon).

All of the same considerations in setting playbook values described in Link KPIs to playbooks and run them manually, on page 75 apply to this automatic option. Note, however, that:

- You must ensure that none of the required linking parameters are left empty. The user interface indicates the required parameters.
- You must not set any of the form fields as "runtime" parameters. If you are running playbooks automatically, you will not have the option to choose a value at runtime.
- If you are a non-admin user, ensure that you have access to the **Auto Remediation** task. Unless you have access to this task, you cannot unlink or link KPIs to playbook with automatic remediation.

  Prerequisites:

  - Ensure that the Health Insights application is installed.
  - Ensure that **Playbook job scheduling** is enabled and **Credential prompt** is disabled in the Device Override Credentials page. For more information, see Verify installation and configure system settings, on page 5.

  You must have Crosswork system administrator privileges to change these settings. Once these settings are saved, you cannot change them unless you first use the Crosswork Manager to uninstall, then reinstall both the Change Automation and Health Insights applications.

  To enable the task permission, do the following:

  1. Go to **Administration** > **Users and Roles** > **Roles**.
  2. Under the **Roles** pane, select the role for which you want to grant the access.
  3. Under the **Task Permissions** tab, enable the **Auto Remediation** check box and click **Save**.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Key Performance Indicators (KPI)**. The **Key Performance Indicators (KPI)** window opens, listing the KPI categories and the KPIs available in each category.

Step 2 Select the KPI you want to link to one or more playbooks. You can use filters to find the KPI you want, as explained in Manage KPIs, on page 70.

Step 3 Click Link Playbook. The **Link Playbook to KPI** window opens.

Step 4 The left side of the window lists the name of the selected KPI and the playbooks appropriate for linking to it. Scroll through the list, or use the **Playbook name** field to restrict the list to just the playbooks you want to see.

Step 5 When you have found the playbook that you want to link to your chosen KPI, click the playbook name. The right side of the window will then list the **Playbook details** for the selected playbook, including:

- The **Hardware platform** and **Software platform** with which the playbook is compatible.
- The minimum software **Version** required to execute the playbook.
- The **KPI alert severity** level needed to trigger a run of this playbook. Note that, if you will be choosing multiple playbooks to run when a KPI alert is raised, be aware that playbooks do not share severity levels. If you have selected a **Critical** severity level for one playbook, you must select **Major**, **Minor**, **Warning** or **Info** severities for a second playbook, and another still for a third playbook.
- Under the **Set playbook execution** field, select **Automatic**. Note that, if you or a Crosswork administrator have not already done so, Crosswork will prompt you to enable **Playbook job scheduling** (and disable **Credential prompt** overrides) in order to enable automatic playbook execution.
- Modify the **Set playbook parameters** default values to be used when the playbook runs. In many cases, you can select from a range of default values. You can also enter your own. These values vary widely, depending on the playbook and its purpose. For help, see the information offered on screen for the playbook you have selected.

Step 6 Verify or modify the **Source** and other parameter values as needed.

Step 7 When you are finished making changes, click **`Link to KPI`**. The **Key Performance Indicators (KPI)** window is displayed again, this time with the linked playbooks shown next to the name of the KPI in the **Key Performance Indicators (KPIs)** list.

Step 8 If you want to run more playbooks (up to three playbooks total): Repeat steps 5 through 7 for each additional playbook you want to run when an alert for this KPI is raised.

Step 9 To change the playbook parameters linked to a given KPI, repeat steps 5 through 7 for that KPI, but this time choose the playbook whose settings you want to modify. If you have chosen multiple KPIs, you can switch among them by clicking on the playbook tiles at the top of the window. To unlink a playbook entirely, select the KPI and click Unlink.

Manage KPI profiles

The Health Insights KPI profiles window allows you to create, edit, and delete KPI profiles.

A KPI profile is a collection of KPIs and their corresponding parameters such as alert frequency, alert type, cadence, and more. You can group relevant KPIs into a KPI profile, give it meaningful name that is based on

the purpose (for example, environmental or health check), and configure parameters that are relevant to monitoring a specific type of devices (for example, edge routers). Once the KPI profiles are created and validated by the system, they are ready to be used. You can select the device(s) in Health Insights, select appropriate KPI profiles, and enable them. This action enables all the KPIs in the selected KPI profile. Similarly, you can select the device(s) and choose to disable the KPI profiles. This removes all the collection jobs on the Crosswork Data Gateway for all the KPIs and for MDT-based KPIs, this removes the configuration in the device(s).

To display the Health Insights KPI profiles window, choose **Performance Alerts** > **KPI profiles** from the main menu.

Figure 49: KPI profiles

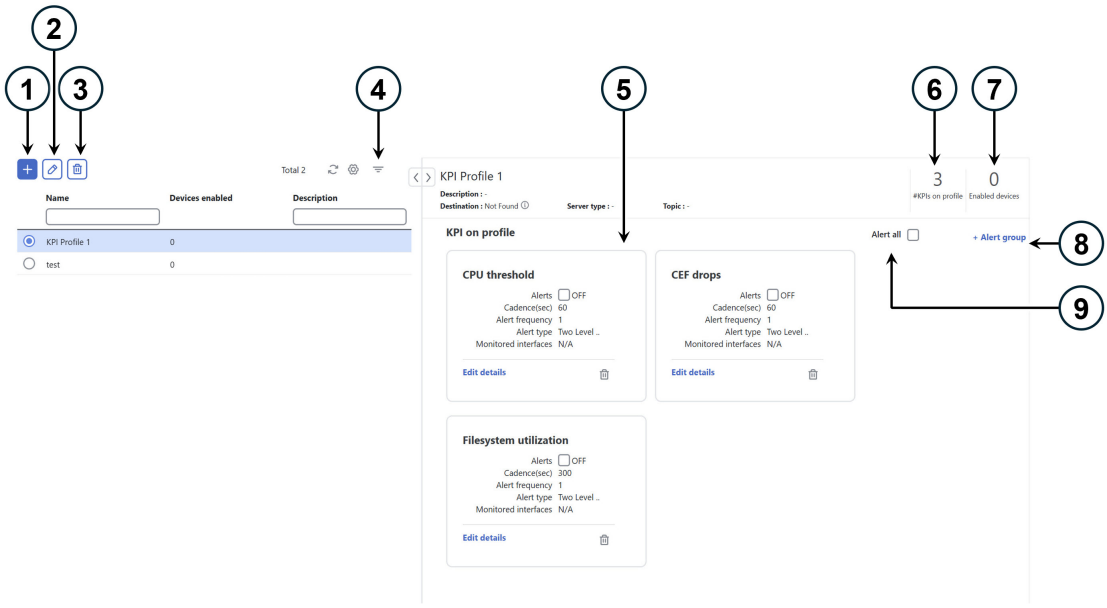

Item	Description
1	**Create KPI profile**: Click + to create a new, user-created KPI profile. For help with this task, see Create a new KPI profile, on page 80.
2	**Edit KPI profile**: Select a user-created KPI profile in the list and then click ✎ to edit it.
3	**Delete KPI profile**: Select a user-created KPI profile in the list and then click 🗑 to delete it. You cannot delete a KPI profile that has been enabled on any device(s).
4	**Filter KPI profile**: To find a KPI category, enter all or part of the KPI profile name in this field, and the list is automatically filtered based on your input. Click ⚙ to clear any filters you have set. Filtering is case-sensitive.
5	**KPI on profile**: The KPI(s) added on the selected KPI profile and the associated parameters are displayed here. You can edit the KPI parameters, or remove a KPI from the selected KPI profile using the appropriate options here. For KPI profiles with custom KPI, the Alert checkbox will be disabled, and an alarm will be raised to inform users that alerting is disabled for the profile.

Item	Description
6	**#KPIs on profile**: This is the number of KPIs added on the selected KPI profile.
7	**Enabled Devices**: This is the number of devices on which the selected KPI profile is enabled.
8	**+Alert Group**: Click this option to create Alert Group for the selected KPI profile. For help with this task, see Create a new KPI profile, on page 80
9	**Alert All**: Click this option to turn off or turn on the alerts for all KPIs in the profile.

Create a new KPI profile

You can create a KPI Profile and enable it on the desired devices. The workflow is as follows:

1. Supply basic information, such as the profile name and a description.
2. Add KPI(s) and save the profile.
3. Edit KPI parameters and create alert groups.
4. Enable the KPI Profile on the devices.

The following steps explain how to perform these tasks.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **KPI Profiles**. The **KPI Profiles** window is displayed.

Step 2 Click the  . The **Create New Profile** window is displayed.

Step 3 In the text fields provided, enter a unique **Profile name**, a short **Description**. The **Profile name** can contain a maximum of 32 alphanumeric characters, plus underscores ("_"). No other special characters are allowed.

To avoid problems with alerting, ensure that each KPI **Profile name** you assign is unique and does not share sub strings with other KPI profiles. For example: In a set of three KPI profiles with the IDs "L2", "L2SNMP" and "L2GRPC", all three profiles IDs contain the sub string "L2".

Step 4 (Optional) You can specify an external destination to send the data collected by KPIs. To create an external data destination, go to **Administration** > **Data Gateway Global Settings** Provide relevant values for the following fields:

- **Server type**: Select either KAFKA or GRPC.
- **Name**: Select the name of the external destination.
- **Topic**: Enter a topic to provide context for the data being sent. This field is applicable only for KAFKA.

Note
You need to create a new data destination to export the KPI data. The predefined data destinations cannot be used for this activity. For more information about creating a data destination, see the Add or edit a data destination in the *Cisco Crosswork Network Controller Administration* guide.

Step 5 Add KPI to the profile, using the following filter options:

a) **All KPIs**: By default, this option is selected and all available KPIs are displayed in the list. You can select the desired KPI by checking the relevant check box.
b) **Recommended KPIs**: Use this option if you want the system to recommend KPIs for your devices. Clicking **Recommended KPIs** displays a list of your devices from the network. You can filter the device list by entering relevant values in the Name field, or by using tags. Select a device from the list to show recommended KPIs on the right side. Select the desired KPI by checking the relevant check box.

Note
Selecting KPIs from the recommended KPI list of a selected device does not automatically enable the KPI Profile in the selected device. The KPI Profile can be enabled after it is created. For more information, see Enable KPI profiles on devices, on page 83.

Step 6 Click **Save** save the new KPI profile and display the **KPI profiles** window.

Step 7 In the **KPI profiles** area on the left side, choose the KPI profile that you created, and the individual KPI details are displayed on the right side.

Note
For the Interface KPIs, you can gather the data for **all** the interfaces or **selected** interfaces. If you opt to gather the information for **all** the interfaces, a warning symbol appears on the KPI profile name on the left side and on the individual KPI details on the right side, indicating that the monitoring interfaces are not customized.

Important
Gathering telemetry data for all the interfaces can be resource-intensive and may require additional worker nodes and/or CDG resources to be deployed.

Step 8 You can leave the KPI parameters at the default or choose a different value. To edit the KPI parameters and preferences, click **Edit details**, and the **KPI details** window is displayed. Edit the values as appropriate for the purpose of your KPI. The details are:

- **Common parameters**
  - **Alert**: This is an on/off toggle switch for alerting. Based on the **Alert** parameter value, the corresponding alerting logic is deployed. Alerting can be enabled even after the KPI profile has been applied to the devices.

    Note
    Any KPI using the group alerting logic need to have the alerting flag set to ON.

  - **Cadence(sec)**: Set the frequency of sensor data. Set the frequency (in seconds) in which the KPI will gather sensor data from the devices on which the KPI profile is enabled.
  - **Alerting down sample rate**: Alert frequency rate. It determines how often KPI data will be evaluated for any alert conditions, and is relative to the Cadence. For example, if Cadence is 60 seconds and you want to do an alerting evaluation every 300 sec, then specify Alerting Down Sample Rate as "5".
- **KPI monitoring preferences**: Applicable only for Interface KPIs.

Figure 50: KPI monitoring preferences

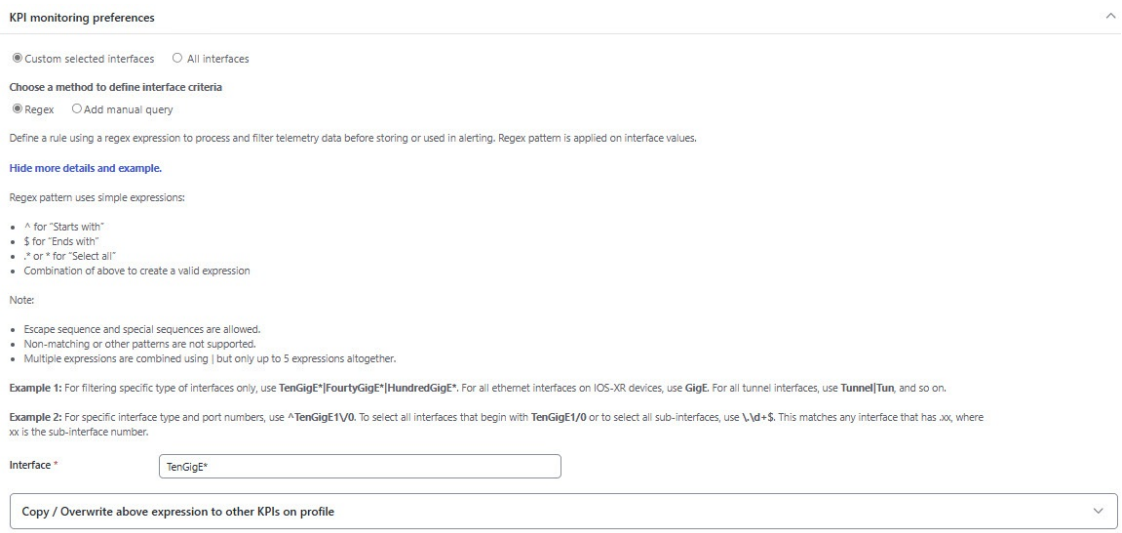

- **Customer selected interfaces**: You can define the interface criteria.
  - **Regex**: You can define a rule using regex expression.
  - **Add manual query**: You can add different sets of rules.
- **All interfaces**: The selected KPI is applied to all the interfaces.

Step 9 You can also edit the alert logic parameters of the selected KPI. To learn more about a parameter, hover your mouse cursor over the shown next to the parameter name.

Note
When different thresholds are desired for different types of devices in the network, it is advisable to create multiple profiles and split the KPIs across them to meet the needs of different device types.

Step 10 When you are finished making changes, click **Save** to save the new KPI profile. Health Insights validates your input parameters and displays the **KPI Profiles** window.

Note
You can create up to 50 KPI profiles, and an individual KPI profile can consist up to 50 KPIs. KPI profile creation can fail if the total number is exceeded, or if Health Insights could not create the required tags in Inventory manager. This status is reflected in the profile state. Once profile is ready, it can be applied on devices.

With the **KPI Profiles** window displayed, you can enable the new KPI profiles on one or more devices immediately, following the steps given in Enable KPI profiles on devices, on page 83.

See Disable KPI profiles on devices or device groups, on page 86 for instructions to disable KPI profiles.

Step 11 (Optional) You can also create alert groups for a KPI profile. Alert groups use Boolean logic (cascaded OR and AND) to combine alert outputs from primary KPIs in your KPI profile and create a group logic query. To create an alert group, click **+ Alert group**. The **Create alert group** window is displayed.

Note

Configuring an alert provider enables the alerts from the group alert to be sent to a REST endpoint using Webhook registered in the alert provider.

Step 12 Provide a relevant entry in the **Name** field. **Summary** and **Details** are optional fields.

Step 13 The **Alert group conditions** area on the right side lets you select a logic gate (AND/OR) and add a KPI on which the logic is applied. Your alert group can be based on the alert criteria of a single KPI, or it can be a combination of multiple KPI outputs. Click the desired logic (**AND** gate is selected by default), and click the **+ ADD** dropdown list to add an **Item** or a **Group**.

Item allows you to add individual KPI items and set the corresponding alert level, and **Group** allows you to add a nested alert group.

Step 14 Choose the desired KPI from the **Select KPI** dropdown, and select the desired level(s) for which the alerts need to be set for the chosen KPI. The alert levels are CRITICAL, MAJOR, MINOR, WARNING and INFO. Based on the logic gate and alert criteria you select, the output of the KPIs are evaluated and the alert is generated.

Figure 51: Create alert group

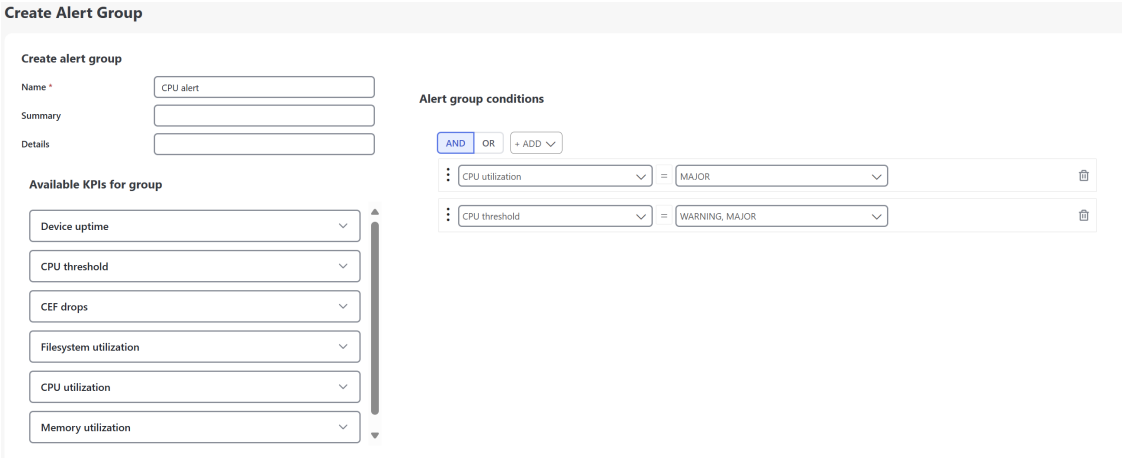

In the example shown above, the alert is set based on the output of two logic gates. The first logic gate is the output of an **OR** operation between the **Memory utilization** and **Interface bandwidth monitor** KPIs. If the set alert levels are met for either of the KPIs, the output of the first logic gate is set as true. This output is considered as the input for the second logic gate, which is an **AND** operation with the **CPU utilization** KPI. If the alert levels of both the KPIs are met, the output of the second logic gate is set as true.

Step 15 Click **Save** to save the new alert group and display the **KPI Profiles** window. Click **Edit details** or to edit or delete an existing alert group respectively.

Enable KPI profiles on devices

With Health Insights, you can enable and monitor the KPI profiles in which you are interested. Instead of sifting through all the data that a given device can supply, you choose to monitor only the information relevant to the role the device plays in your network. Your network devices operate most efficiently when configured to only report data that specifically relates to the performance of its role in the network.

Some KPIs trigger alerts based on deviation from an established level of performance. For these types of KPIs, it is necessary to allow the system some annealing time in order to establish normal performance levels.

Important

- You can only enable KPI profiles with MDT-based KPIs on a device that has been mapped to a Cisco Network Services Orchestrator (Cisco NSO) provider and attached to a Crosswork Data Gateway.
- Do not enable KPI profiles on devices that are not reachable.
- The load that is created on Data Gateway and Platform Infrastructure caused by enabling KPI profiles on many devices or KPI profiles that gather a lot of data is hard to estimate. Crosswork provides a UI and API that allows you to see the current load and provides general guidelines for determining when you must refrain from enabling more collections until either other collections are disabled or more resources (Data Gateway or worker nodes) are added. To check Data Gateway load, see Crosswork Data Gateway health metrics in the *Cisco Crosswork Network Controller Administration* guide.

To enable KPI profiles on devices:

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Enable/Disable KPI Profiles**. The **Devices** window is displayed.

Step 2 Select the devices for which you want to enable KPI profiles. You can click the **Device** or **Device tags** buttons above the table on the left to toggle between selecting the devices by name or by tagged device group membership. Depending on your selection, the device list or the device tag list is displayed on the left.

Figure 52: Devices window

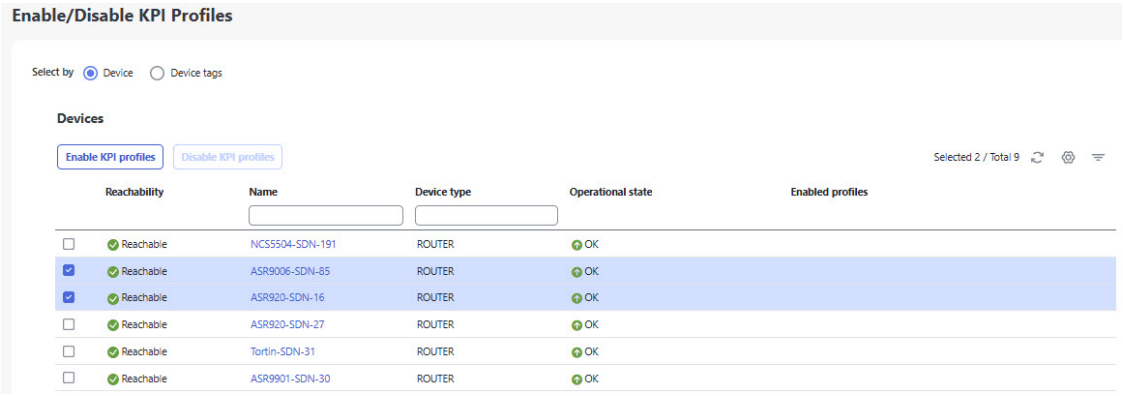

If you choose to select by **Device**:

- Click in the table on the right. Type a **Name** or **Device type** in the filter fields. As you type, the table displays only the devices whose name or type match the text that you typed.
- Click the check box next to the device(s) you want. You can select multiple devices at the same time.

If you choose to select by **Device tags**:

- Type a tag name in the **Name** field to find a Device Group in the table. As you type, the table displays only the tag names that match the text you typed.
- Click the check box next to the group that you want. The names of all the devices in that group appear in the devices table on the right.

Note
You cannot enable a KPI on a device that is attached to a standard Crosswork Data Gateway. Also, you cannot move a KPI-enabled device from extended Crosswork Data Gateway to standard Crosswork Data Gateway. In both cases, Crosswork displays an error pop-up.

Step 3 Click **Enable KPI profiles** to continue. Health Insights detects the selected devices, their types and models, and retrieves and analyzes their running configurations. The **KPI profiles** window presents the KPI profiles available for your selected devices.

Step 4 Choose the KPI profiles that you want to enable by clicking the check box next to the KPI profile name, and click **Next**.

The **Verify details** window appears, listing all the KPI profiles you have chosen to be enabled on the selected devices.

Step 5 (Optional) To get information about the KPIs included in the KPI profile. Click the KPI profile in the **Selected profile(s)** table, and the content of the selected KPI profile is displayed on the right side. Click **View more details** to view the parameters of a specific KPI. A pop-up window provides the details of the KPI. Click the ✕ to close the pop-up window.

Step 6 To enable the selected KPI profiles on the selected devices, click **Enable**. Health Insights schedules the KPI profile(s) as a series of job sets.

The **Alert** flag for the KPI profile (click **Edit details** on the relevant KPI) must be turned **ON** in order to trigger an alert when the data is collected.

Enabling a KPI results in configuring a collection job on the Crosswork Data Gateway. For GNMI-based and SNMP-based KPIs, the Crosswork Data Gateway polls the desired data and forwards it to Health Insights for processing and evaluation. For MDT-based KPIs, the devices (through NSO) are configured to push the data to the Crosswork Data Gateway which then forwards it to Health Insights for processing and evaluation.

In the **Device** table, in the **Enabled profiles** column, you can click the number to see the status of the KPI collection job (for example to see if the KPI profile ID is active or not).

Step 7 From the main menu, choose **Performance Alerts** > **KPI Job History** to watch the progress of each job set, as shown below. You should see job sets completing with a status of "Success". If the job sets complete with a "Partial" or "Failed" status, ensure to read the job completion messages, and check that the selected devices are still reachable.

Figure 53: KPI Job History

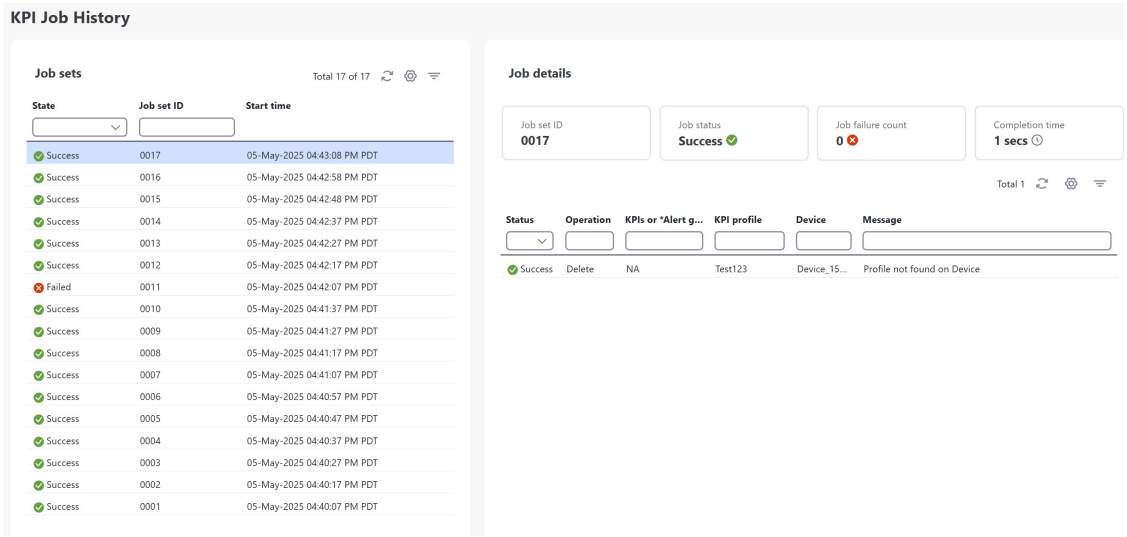

When the job sets complete successfully, the KPIs are now associated to the devices and the platform begins the process of enabling the relevant collection procedures for those network elements. In making these changes, you are automating the configuration of both the platform and the devices themselves to collect only the information required.

Step 8 From the main menu, choose **Administration** > **Collection Jobs** to look at the collection jobs and to make sure that they are created and the incoming data is collected. For more information on monitoring the status of collection job, see the Monitor collection jobs section in the *Cisco Crosswork Network Controller Administration* guide.

Step 9 From the main menu, choose **Performance Alerts** > **Alert Dashboard**. The Health Insights alert dashboard, on page 87 shows the alert status for the devices on which you have enabled KPI monitoring.

- SNMP/MDT jobs may take more time than expected to reach the completed state when there is an increase in the number of devices, interfaces and KPIs.
- Enabling KPI profiles per device takes around 3-5 seconds (but the time varies based on the number of KPIs being enabled). If the device is not reachable, it keeps trying until it is timed out. This may result in the job taking more time to reach the completed state.

Verify the deployment status of enabled KPIs

After you enable a KPI profile, you can verify the deployment status.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **KPI Job History** The **KPI Job History** window lists the jobs that have been run most recently, indicating whether they succeeded or failed, when they ran, and on what devices.

Step 2 Click the transaction ID in the job listing to view detailed KPI job information, including the device on which the KPI profile was enabled and the KPI ID.

Any KPI job stuck in the processing state that does not complete within 60 minutes will be marked as "failed". After addressing any underlying issues (for example, device connectivity, credentials, NSO sync, and so on), the same job must be reactivated, as explained in Create a new KPI, on page 73.

Disable KPI profiles on devices or device groups

You can use the **Enable/Disable KPI Profiles** window to disable the KPI profiles running on device(s) or device groups.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Enable/Disable KPI Profiles**. The **Enable/Disable KPI Profiles** window is displayed.

Step 2 To disable KPIs enabled on one or more devices:

a) Click the **Device** button above the table on the left. The **Devices** table displays all the devices, with the total number of KPIs enabled on each device.
b) Click the check box next to the devices on which you want to disable KPIs.

   If you select one device, you can disable all KPI profiles for the device or just some of the KPI profiles. If you select more than one device, you can only disable all KPIs for them.

c) Click **Disable KPI profiles**. You will be prompted to confirm that you want to disable the KPIs running on all the selected devices. If you selected only one device, click the checkboxes next to the KPI profiles you want to disable on that device, or click the checkbox at the top of the column to disable all the KPI profiles running on that device. Click **Disable** to confirm.

Step 3 To disable all KPI profiles enabled on all the devices within a device group:

a) Click the **Device tags** button above the table on the left. The table displays the list of device tags.
b) Click the checkbox next to the device tag(s) used on the devices from which you no longer want to collect the KPI data.

   When you select a device tag, the **Devices** table on the right shows all the devices that are associated with that tag. All of the devices are preselected.

c) Click **Disable KPI profiles**. You will be prompted to confirm that you want to disable all the KPIs running on all the devices in the group. Click **Disable** to confirm.

Health Insights alert dashboard

The Health Insights alert dashboard provides device health summary information that is based on real-time network state events. The dashboard displays a network view of KPI sensors that are paired to specific device groups. Health Insights raises customizable events and alerts that are based on user-defined logic.

Note Alert dashboard displays individual KPI alerts, although the mechanism of enabling KPI on a device is done through a KPI profile.

To display the Health Insights dashboard, choose **Performance Alerts** > **Alert Dashboard** from the main menu.

Figure 54: Health Insights Alert Dashboard

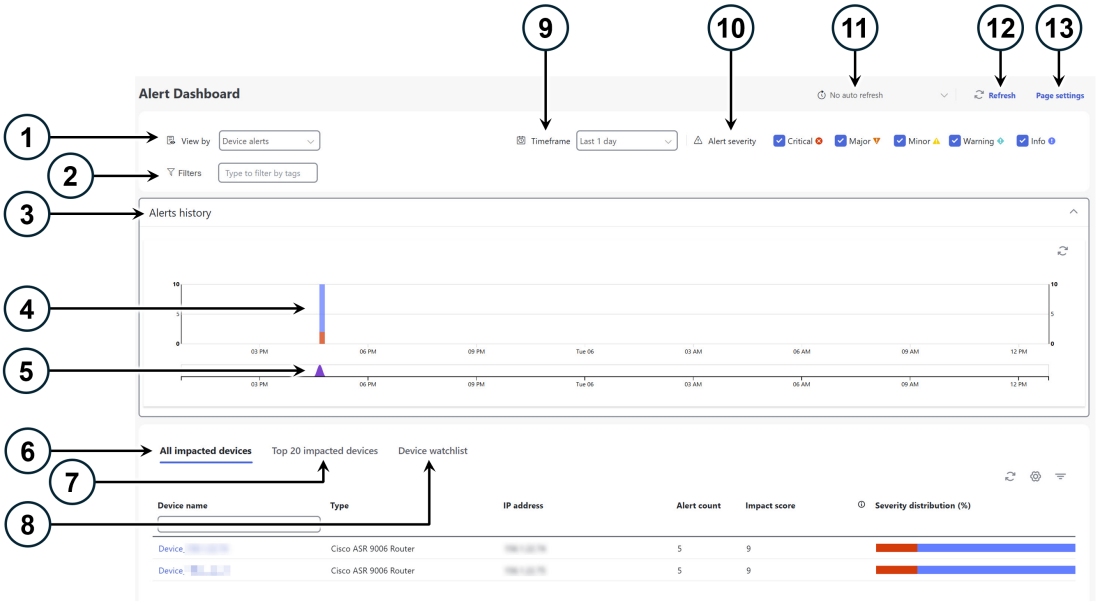

Item	Description
1	**Device/KPI Alert selector**: Click here to toggle between device alert and KPI alert information.
2	**Filters**: This field lets you filter the alert dashboard information by associated tag names. To select a tag, do one of the following: • If you know the tag that you want to use, enter it in the **Type to filter by tags** field and then check its check box. Repeat this step to select more tags. • If you want to select a tag from the tags that are currently available: 1. In the **Type to filter by tags** field, type any character to open the results list. 2. Click the **View all tags** link at the bottom of the list. 3. Check the check box for each tag that you want to use and then click **Apply filters**. 4. Delete the character that you typed in Step 1 to clear the results list. Tag filters you create are not saved. If you open another window and then return to the alert dashboard, you need to re-create tag filters.
3	**Alerts history**: This dashlet shows the total number of device alerts or KPI alerts that have been raised during the chosen time period, with detailed time lines showing both individual sets of alerts and the overall alert trend.
4	**Alerts history bar graph**: This graph shows alerts as discrete bar indicators whose height represents the total number of alerts gathered at each point in time. To see the total for each type of alert, hover your mouse cursor over the bar indicator. You can also use the **Alerts trend** line to zoom in on particular portions of the alert history.

Item	Description
5	**Alerts trend line**: This line shows the overall trend in alerts for the chosen time period. You can use the **Alerts trend line** to select and zoom in on a specific time period within the **Alerts history line**, as follows: **1.** Click the time-period starting point in the **Alerts trend line** and hold down the mouse. **2.** Drag the cursor to the endpoint and then release the mouse. To restore the full view of the **Alerts history line**, click on any point outside of the light gray shading on the **Alerts trend line**.
6	**All impacted devices/All impacted KPIs**: When selected, this dashlet provides a complete list of all devices or KPIs affected by alerts. The information for each affected device or KPI includes: • Device name or KPI name • Device or KPI type • IP address: The IP address of the impacted device. This column is only displayed for devices. • Alert count: The total number of alerts for that device or KPI during the selected period. • Impact score—This value is determined using the following formula: (5 x number of critical alerts) + (4 x number of major alerts) + (3 x number of minor alerts) + (2 x number of warning alerts) + (1 x number of info). These are the default values. You can change the weightage in the **Page settings** option. When monitoring the health of your network, focus on devices or KPIs with a higher impact score. • Severity distribution—Provides a visual breakdown of the severity that is associated with a device or KPI's alerts. To view a tooltip that indicates the number of raised alarms (by severity and in total), place your cursor over the appropriate bar segment.
7	**Top 20 impacted devices/ Top 20 impacted KPIs**: When selected, this dashlet displays a map of tiles, each tile representing one of the 20 devices or KPIs with the most alerts during the selected time period. The amount of space that each tile occupies in the map corresponds to the number of alerts raised: the more alerts, the bigger the tile. Also, the tiles are color coded. The colors correspond to the **Alert severity**. To view more detailed information for a particular device or KPI, click the device or KPI name link in the center of the tile.

Item	Description
8	**Device/KPI watchlist**: When selected, this dashlet provides a list of all devices or KPIs, that you had selected from **+ Manage device/KPI watchlist**, which are affected by alerts. The information for each affected device or KPI includes: • Device name or KPI name • Device or KPI type • IP address: The IP address of the impacted device. This column is only displayed for devices. • Alert count: The total number of alerts for that device or KPI during the selected period. • Impact score—This value is determined using the following formula: (4 x number of critical alerts) + (3 x number of major alerts) + (2 x number of minor alerts) + number of warning alerts. When monitoring the health of your network, focus on devices or KPIs with a higher impact score. • Severity distribution—Provides a visual breakdown of the severity that is associated with a device or KPI's alerts. To view a tooltip that indicates the number of raised alarms (by severity and in total), place your cursor over the appropriate bar segment.
9	**Timeframe**: Specifies the time period for which the dashboard provides alert information: The last one hour, last day, last three days, last week, and last month. Please note that the dashboard provides alert information only, not telemetry information.
10	**Alert severity**: Maps the bar indicator colors that are used in the **Alert History** dashlet to the corresponding alert severity. To display or hide the alerts for a particular severity, click the check box for that severity. An enabled check box indicates that alerts of that severity have been raised and are being displayed. A clear check box indicates that the alerts of that severity are either not being displayed or have not been raised during the displayed time period.
11	**Auto refresh**: Specifies how often the dashboard is automatically refreshed.
12	**Refresh icon**: Refreshes the dashboard.
13	**Page settings**: Provides the default page settings for that particular session. You can customize the page display based on Alert type, Timeframe, Auto refresh, Detail display, and Alert severity. You can also change the weightage here for the impact score calculation.

Note The individual alerts for any specific KPI are shown in the dashboard. Alerts resulting from the alert group logic are not shown in the dashboard. Only the API shows the impacted results.

View alerts for network devices

After enabling KPIs on a device, you can view alerts for that device and get data for each performance indicator being monitored.

Note The KPIs shown in the following steps are examples. There are many more KPIs available in Health Insights. For the complete list, see List of Health Insights KPIs, on page 64.

Procedure

Step 1 From the main menu, choose **Performance Alerts** > **Alert Dashboard**. The Health Insights Alert dashboard is displayed.

Figure 55: Alert Dashboard

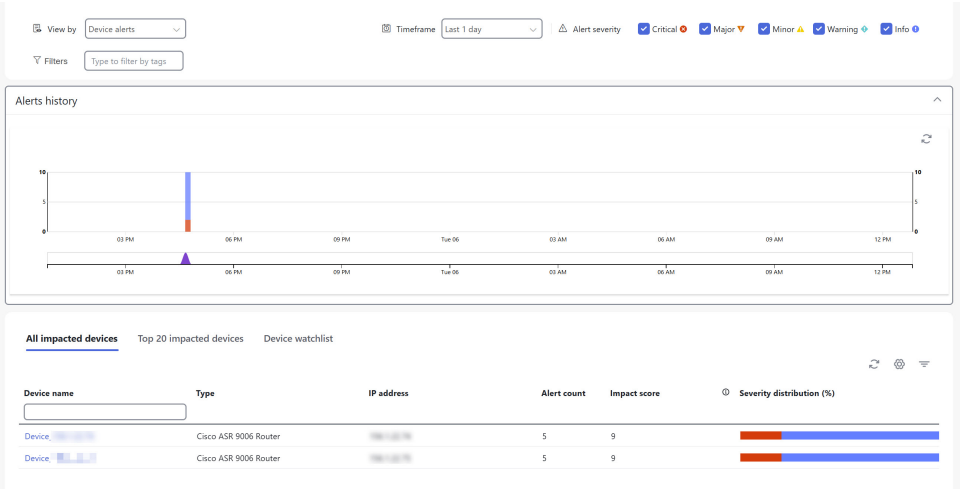

Step 2 Make sure that the **Device alerts** view is displayed (select the **View by: Device alerts** toggle, if needed). Then scroll down below the **Alert history** panel and click the **All impacted devices** tab. The dashboard displays a list of devices with alerts.

Step 3 Click the **Device name** for the device whose details that you want to view. Health Insights displays the device's basic **Overview** information, **Alert history**, a **Topology** map, and the list of the device's currently **Enabled KPIs**.

Figure 56: Device alert details

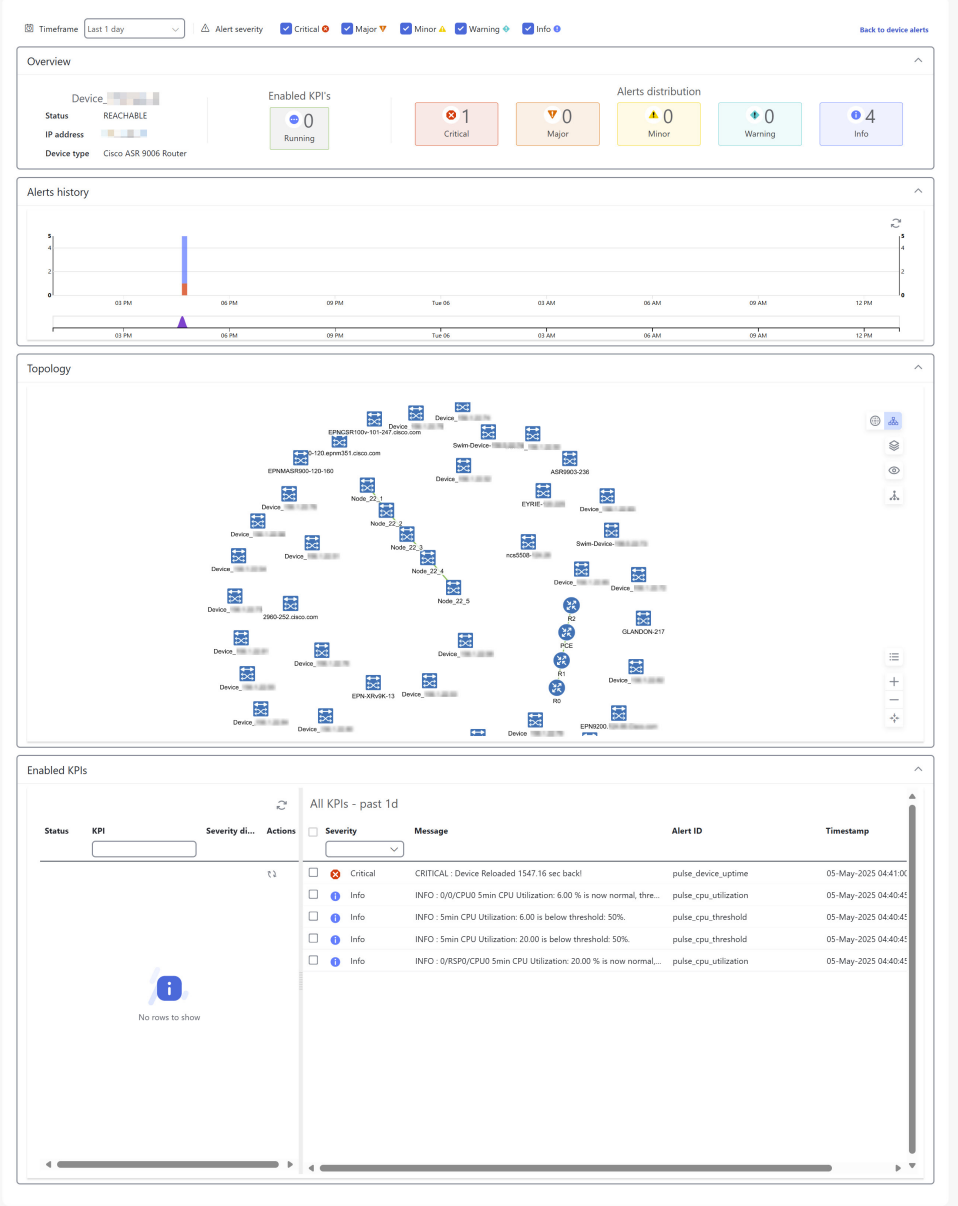

The **Topology** map is a version of the map that you see when you select **Topology** from the main menu, but centered on the device for which you are viewing KPI alerts. The **Enabled KPIs** panel lists all the KPIs currently enabled on the selected device, plus a list of all the alerts for that device raised by any of the enabled KPIs during the past hour.

To see alerts for a different period, click the **Timeframe** dropdown (shown below) and select the time frame you want (up to **`Last 1 Month`**).

Figure 57: Timeframe dropdown

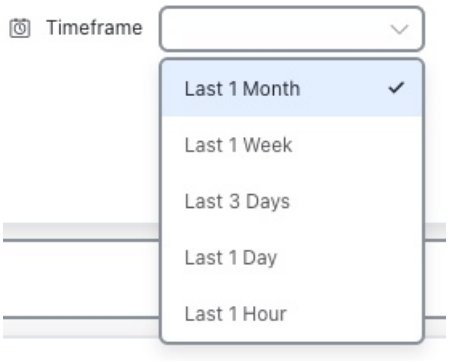

To focus the display only on alerts of the severity you want, check or uncheck the boxes in the **Alert severity** field, (shown below).

Figure 58: Alert severity

Step 4 To view telemetry data received for any of the KPIs for this device: In the **Enabled KPIs** list on the left, click the icon next to the KPI whose telemetry data you want to see. Crosswork displays a popup telemetry data window like the one shown below. The popup window shows a timeline at the top, representing all the alert data received during the last 72 hours (with hourly slots), and relevant performance for the same period in a Grafana graph at the bottom.

Step 5 The timeline shows a blue box, with brushes on the sides, representing the limits of the time period shown in the graph at the bottom. Click on and move the blue box or the brushes on the timeline to select the desired time slot (up to 6 hours). Move the mouse cursor over any data point in the graph to view additional pop-up information for that data point.

A red line or tag represents a point at which the KPI was triggered. This can occur on any subscribed statistic the KPI is monitoring. Health Insights collects and identifies the time points and frequency, which help determine when these events become an operational concern.

Graphical data is only visible for time slots for which alerts were triggered. Be default, the Grafana graph shows telemetry for the last six hours.

Step 6 To focus the Grafana view on a different timeframe, click the time period field (with the clock icon) shown at the top of the **Summary** tab. You can select time periods up to several years.

Telemetry data retention in the Alerts dashboard

Telemetry data collected from devices is retained in the time-series database for the last 72 hours. This data is used in the Health Insights Alert dashboard to identify alerts using a process known as stream-based alerting. The resulting 'alerts,' if any, are stored in the same database and are retained for 30 days. The messages showing the alerts' duration are displayed at the top of the Device/KPI view in the Alert dashboard. For more information, see View alerts for network devices, on page 90. The alerts can also be queried using REST APIs. For more information, see the Cisco Crosswork Network Controller API Documentation on Cisco DevNet.

Troubleshoot Health Insights

The following table describes issues that you may encounter when using the Health Insights application, and their solutions or workarounds.

Table 3: Health Insights Troubleshooting

Issue	Solution
Apply a KPI to a device fails with messages indicating that Cisco Network Services Orchestrator (Cisco NSO) and the target device are out of sync or otherwise out of communication. Message text will vary, but may include "device out of sync," "NC client timeout," and other text indicating that there are connectivity or sync issues between NSO and the device.	1. Go to **Performance Alerts** > **KPI Job History** and check the **Message** column for an error message. 2. Go to **Device Management** > **Network Devices**. 3. For the failed device, in the **NSO state** column, click ⓘ. 4. From the **Check sync** drop-down list, click **Sync from**. 5. Confirm that the device is in sync now.
Operation timeouts can occur when adding a new KPI to an existing KPI Profile and then editing the newly added KPI.	Write times for KPI edits can take up to five minutes, so the edited KPI in the profile will be enabled eventually. If you find the timeout message a problem, you may want to disable the KPI profile for a short period until the write delay has passed.
Health Insights not receiving data.	1. Confirm that the KPI configuration job completed without error: Go to **Performance Alerts** > **KPI Job History** 2. Check the Collection/distribution status: Go to **Administration** > **Collection Jobs**. 3. Check for the collection job to see if the table (accessed by clicking the graph icon for the job) indicates that data collections are processing.