

Juniper Data Center Assurance User Guide

Published
2026-06-10

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Juniper Data Center Assurance User Guide
Copyright © 2026 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | ix

1

Get Started

Juniper Data Center Assurance Overview | 2

User Activation and Login | 3

Reset Your Local Account Password | 7

2

Organization Management

Organization and Sites Overview | 9

Add an Organization | 9

Delete an Organization | 10

Manage Organization Settings | 11

Authentication Methods Overview | 15

Manage Identity Providers | 16

Add an Identity Provider | 16

Edit an Identity Provider | 18

Delete an Identity Provider | 18

Field Descriptions | 18

Manage Roles | 20

Add a User-Defined Role | 21

Edit a User-Defined Role | 21

Delete a User-Defined Role | 22

Manage API Tokens | 22

Add an API Token | 23

Edit an API Token | 24

| Delete an API Token | 24

Integrate Your Juniper Support Resources to Your Organization | 24

Set a Preferred Portal for an Organization | 27

3

Site Management

About the Sites Page | 30

Edit Site Configuration | 31

4

User Management

About the Administrators Page | 37

Predefined User Roles Overview | 38

Add Users to an Organization | 40

Invite Users | 42

Manage Users and Invites | 44

| Edit User Role | 44

| Reinvite a User | 45

| Cancel an Invitation | 45

| Revoke a User | 46

Manage Your Account | 46

5

Inventory Management

About the Inventory Page | 51

DC Edges and DC Switches | 64

Validate Your Juniper Account Overview | 67

Validate Your Juniper Account | 69

| Validate Your Juniper Account | 71

| Validate Your Juniper Account as an Unregistered User | 72

| Users with Validated Juniper Account | 73

[Configure Ports | 75](#)

6

Contracts Management

[About the Contracts Page | 80](#)

7

Advisories

[About the Advisories Page | 90](#)

8

Support Insights

[About the Support Insights Page | 99](#)

9

Audit Logs

[Audit Logs Overview | 105](#)

[About the Audit Logs Page | 105](#)

10

Licensing

[Juniper Data Center Assurance Licensing and Subscriptions | 109](#)

11

Marvis AI Assistant for Data Center

[Marvis AI Assistant for Data Center | 113](#)

[About the Marvis Page | 113](#)

[View Data Center Events in Marvis Actions | 114](#)

[Event Types Displayed in Marvis Actions | 115](#)

[Marvis AI Assistant Overview | 118](#)

[Find Information and Troubleshoot Juniper Products and Services Using Marvis AI Assistant | 119](#)

12

Monitor and Troubleshoot Data Center Events

[Adopt Juniper Apstra Edge | 124](#)

[Apstra Edge Overview | 124](#)

[Adopt Apstra Edge in Juniper Data Center Assurance | 124](#)

[Verify DC Assurance Setup | 128](#)

[Upgrade Apstra Edge from Juniper Data Center Assurance | 134](#)

About the Dashboard Page | 136

Data Center Network Topologies on DC Assurance | 148

Access Apstra Data Center Director from Juniper Data Center Assurance | 158

Monitor and Troubleshoot Data Center Events from Mist | 160

13

Application Awareness

Application Awareness Overview | 166

About the App Aware Page | 171

About the Service Topology View Page | 173

14

Assurance

Impact Analysis Overview | 176

Impact Analysis Overview | 176

Simulate Failure Scenarios | 178

View Device Details | 178

Benefits of Impact Analysis | 179

About the Impact Analysis Page | 180

Key Anomalies in Impact Analysis | 181

Predictions Overview | 185

About the Predictions Page | 197

15

Network Sustainability Analytics

Network Sustainability Analytics Overview | 201

Forecast Power Cost and GHG Emission | 203

Sustainability: Supported Platforms and Software Versions | 205

Getting Started with Network Sustainability Analytics | 207

Configure Sustainability Settings for the Organization | 213

View Sustainability Insights for an Organization | 216

Understanding the Organization-Level Sustainability Dashboard | 217

Overview | 217

Potential Energy Savings | 225

Configure Sustainability Settings for a Site | 227

View Sustainability Insights for a Site | 229

Understanding the Site-Level Sustainability Dashboard | 231

About the Overview Tab | 232

About the Topology Tab | 243

About the Potential Energy Savings Tab | 249

Review the Sustainability Report | 251

View Sustainability Insights for a Device | 256

Understanding the Device-Level Sustainability Dashboard | 258

About the Overview Tab | 258

About the Power Supply Tab | 268

About the Thermal Profile Tab | 271

About the Recommendation Tab | 276

Configure Port Threshold for the Organization | 278

Configure Port Threshold for a Site | 282

Optimize Energy Consumption by Shutting Down Unused Ports | 285

Optimize Low Traffic Ports Based on Historical Traffic Patterns | 289

Optimize Low Traffic Ports Based on Forecasted Traffic Patterns | 293

Optimize Energy Consumption by Shutting Down Unused Line Cards | 298

Optimize Energy Consumption by Shutting Down Unused PFEs | 303

Network Sustainability Analytics – Troubleshooting | 307

DC Hub

DC Hub Overview | 316

About the DC Hub Page | 321

17

Service Level Expectations

Service Level Expectations Overview | 324

About the Service Levels Page | 333

18

Alerts

About the Alerts Page | 337

Configure an Alert Template | 338

Configure E-mail Notification for Alerts | 338

19

Webhooks Integration

Configure Webhooks to Receive Event Notifications in Third-Party Applications | 341

Webhooks Source Address | 344

Webhook Delivery Events | 345

View Webhooks Delivery Events | 347

20

Help & Support

Provide Feedback About DC Assurance | 352

About the Support Tickets Page | 354

Create a Support Ticket | 360

Requesting Technical Support | 370

About This Guide

Use this guide to understand the features and the tasks that you can perform from the Juniper Data Center Assurance (DC Assurance) application. This guide provides feature overviews and procedures that help you understand the features and perform the various tasks.

Juniper Data Center Assurance is a SaaS-based Day 2 observability platform for data centers that are managed using Apstra Data Center Director.

Documentation Conventions

For information about the notice icons and text and syntax conventions used in the documentation, see [Documentation Conventions](#).

1

CHAPTER

Get Started

IN THIS CHAPTER

- [Juniper Data Center Assurance Overview | 2](#)
 - [User Activation and Login | 3](#)
 - [Reset Your Local Account Password | 7](#)
-

Juniper Data Center Assurance Overview

Obtaining real-time visibility into the functioning of the data center is critical to providing an unmatched user experience. It is important that administrators have the ability to resolve data center events and anomalies proactively, in the same way as administrators resolve campus network events using Mist.

For this, administrators should have visibility into the operations of the data center in real-time. Juniper Data Center Assurance (DC Assurance) is a SaaS-based Day 2 observability platform for data centers that are managed using Apstra Data Center Director. The DC Assurance application, which can receive, process, and perform root cause analysis of network events from a data center managed by Apstra Data Center Director, enables network administrators to proactively respond to data center events ensuring that users' application experience remains unaffected.

Juniper Data Center Assurance runs as an independent application in the cloud, receiving and analyzing events information from Apstra. You can also integrate Juniper Data Center Assurance with Mist. When integrated with Mist, you can view the total number of data center events in the **Data Center/ Application** category in Marvis Actions along with other event types for the campus and branch networks. You can then access Juniper Data Center Assurance and view more detailed information about those data center events in Marvis Actions by clicking the **Data Center Actions** event type under **Data Center/Application**. This way, network administrators get complete visibility into the operations of the entire enterprise network, comprising the campus, branch, and data center networks. Administrators can resolve network issues proactively, which in turn enables enterprises to improve operational efficiency and reduce operational cost and downtime.

Juniper Data Center Assurance supports the following features:

- **Application Awareness**—Provides visibility into application performance by visualizing the services that are active during a given time period. The visualization also shows how the data center fabric is being utilized by these services. You can select a service and drill down to get detailed information about the traffic flow. For more information, see ["Application Awareness Overview" on page 166](#).
- **Impact Analysis**—Provides a graphical representation of the network infrastructure along with the path that the service traffic takes. The network topology provides information about how nodes in the data center are connected to the clients, how application traffic is handled, and how an anomaly in a node can impact application traffic. This enables faster resolution of network issues that impact end user experience. For more information, see ["Impact Analysis Overview" on page 176](#).
- **Predictive Analytics**—An ML-based feature that analyzes historical data points, predicts future metrics, and recognizes outliers in the forecasted data. The insights from this information helps the network admin prevent potential anomalies, device failures, and optical interface failures in the network. For more information, see ["Predictions Overview" on page 185](#).
- **Network Sustainability Analytics**—Provides real-time visibility into power consumption, GHG emissions, and associated operational costs of the data center network. It forecasts power

requirements, GHG emissions, and power costs, which support sustainability reporting. The recommendations for shutting down unused or underutilized components without impacting traffic help conserve power and reduce emissions. The site-level energy dashboard helps you quickly understand the site's energy profile and identify devices with power anomalies. The sustainability feature thus helps improve operational efficiency and supports sustainability goals of the organization. For more information, see ["Network Sustainability Analytics Overview" on page 201](#).

- **Service Level Expectations (SLEs)**—Sets a benchmark for the performance of a data center network. SLEs help network administrators understand whether the network is performing optimally and enable them to respond proactively to network events or performance issues. For more information, see ["Service Level Expectations Overview" on page 324](#).
- **DC Hub**—Provides a comprehensive platform for the network administrators to discover, install, and manage add-ons in your Apstra environment. These add-ons act as plug-and-play options to add functionality and analytics within Apstra for your specific requirements. For more information, see ["DC Hub Overview" on page 316](#).
- **Marvis Actions**—Utilizes the AIOps capability of Marvis AI Assistant for Data Center to analyze data center events and displays them in the Marvis Actions dashboard. Marvis Actions provides recommended actions that enable administrators to proactively respond to data center events. For more information, see ["Marvis AI Assistant for Data Center" on page 113](#).
- **Marvis AI Assistant**—Enables administrators to quickly search within product documentation for relevant troubleshooting information to resolve the events. Marvis AI Assistant for Data Center also provides a conversational interface, which administrators can use to ask queries in natural language and get relevant responses. With the help of Marvis AI Assistant, you can also try to resolve issues, reduce the need to create support tickets, and accelerate the time to resolution. You also have the option to provide feedback on your experience with the AI-assisted issue resolution process. For more information, see ["Marvis AI Assistant Overview" on page 118](#).



Video: [Juniper Data Center Assurance](#)

User Activation and Login

To log in to Juniper Data Center Assurance (DC Assurance) application, you must create an account in DC Assurance and then, activate the account. After you activate your account, you either create an organization or join an organization through an invite.

DC Assurance initiates user activation when:

- You create an account and an organization and access DC Assurance without an invite.

- The superuser invites you to an organization.

Click the link in the invite and complete the login tasks. Your login procedure depends on whether you are an existing user with a DC Assurance account or a new user without a DC Assurance account.

1. To log in as the first admin user without an invite:


- a. Access the GUI directly at <https://dc.ai.juniper.net>.




NOTE: Juniper Networks recommends that you use the latest version of Chrome, Firefox, or Safari browsers to access DC Assurance.

- b. Click **Create Account**.
- c. Type your first name, last name, e-mail address, and password on the My Account page. The password is case sensitive.
- d. Click **Create Account**.

DC Assurance sends a verification e-mail to activate your account.



- e. Click **Validate me!** in the e-mail body.
The New Account page appears.
 - f. (Optional) Click **View Account** to check your name and e-mail address. Click **Back** to return to the Select an Organization page.
 - g. Click **Create Organization**.
 - h. Type a unique name for your organization and click **Create**.
The New Account page appears.
 - i. Click the organization on the New Account page.
2. To log in as a new user with an invite:
- You receive an invite from an administrator to join an existing organization.



- a. Click **Access *organization-name*** in the e-mail body.
The Invite to Organization page opens in your default browser.
- b. Click **Register to Accept**.
The My Account page appears.

- c. Enter your first name, last name, e-mail address, and configure a password.
The password can contain up to 32 characters, including special characters, based on the password policy of the organization.
 - d. Click **Create Account**.
DC Assurance sends a verification e-mail to activate your account.
 - e. In your confirmation e-mail, click **Validate Me**.
The New Account page opens in your default browser.
 - f. Click the organization for which you received the invite.
You are logged in to the application and can access the selected organization's GUI. The tasks you can perform in this organization depends on your user role. See "[Predefined User Roles Overview](#)" on page 38 for more information.
3. To access an invite as an existing user:
- a. Click **Access *organization-name*** in the e-mail body.
The Invite to Organization page opens in your default browser.
 - b. Click **Sign In to Accept**.
The DC Assurance page appears.
 - c. Enter your username and click **Next**.
The DC Assurance login page appears.
 - d. Enter your password and click **Log In**.
The Invite to Organization page appears.
 - e. Click **Continue**.
The Select an Organization page appears.
 - f. (Optional) You can click **View Account** to verify your account details and click **Back** to return to the Select an Organization page.
 - g. Click the organization for which you received the invite.
You are logged in to the application and can access the selected organization's GUI. The tasks you can perform depends on your role. See "[Predefined User Roles Overview](#)" on page 38 for more information.

RELATED DOCUMENTATION

Reset Your Local Account Password

You can reset your password on the login page in the Juniper Data Center Assurance (DC Assurance) GUI. If you had enabled two factor authentication for your account, it will be disabled when you reset your password. You must re-enable two factor authentication after logging into the GUI using your new password.

To reset your password:

1. On the login page, click **Forgot Your Password?**

The Reset Password page appears.

2. Type your e-mail address in the box and click **Send Reset Link**.

A message confirms that the link to reset password is sent to your e-mail address.

The DC Assurance login page appears.

3. Click **Reset My Password** in the message body of the password recovery e-mail in your inbox.

The Set New Password page appears.

4. Type a new password in the Change Password box and click **Change Password**.

A password must contain eight or more characters that are a combination of upper case and lower case letters, numbers 0-9, and special characters.

The DC Assurance login page appears.

5. Type your e-mail address and click **Next**.

6. Enter your new password and click **Log in**.

The Select an Organization page appears.

7. Select or create an organization.

You are logged into the DC Assurance GUI and can view the dashboard of the organization.

2

CHAPTER

Organization Management

IN THIS CHAPTER


- [Organization and Sites Overview | 9](#)
 - [Add an Organization | 9](#)
 - [Delete an Organization | 10](#)
 - [Manage Organization Settings | 11](#)
 - [Authentication Methods Overview | 15](#)
 - [Manage Identity Providers | 16](#)
 - [Manage Roles | 20](#)
 - [Manage API Tokens | 22](#)
 - [Integrate Your Juniper Support Resources to Your Organization | 24](#)
 - [Set a Preferred Portal for an Organization | 27](#)
-

Organization and Sites Overview

An organization represents a customer (for a service provider) or a branch (for an enterprise). An organization can have multiple sites representing the locations where data centers are installed. While a site can have more than one device, a device can be associated with only one site.

You can group sites based on regions, functions, or other parameters for efficient management of the devices. [Figure 1 on page 9](#) represents the relation between an organization, sites, and site groups. In [Figure 1 on page 9](#), an organization has seven sites and three sites groups (Site Group 1, Site Group 2, and Site Group 3). Site 3 and Site 4 are a part of Site Group 1 and Site Group 3 while Site 7 is part of Site Group 2 and Site Group 3.

Figure 1: Organization, Sites, and Site Groups



Add an Organization

An organization represents the customer (for a service provider) or a branch (for an enterprise).

You can add an organization from the login page when you log in to Juniper Data Center Assurance or by clicking the **Utilities** option in My Account page.

To add an organization:

1. Click **Create Organization** on the login page.
The Create Organization page appears.
2. In the **Organization Name** field, enter a name for the organization.
3. Click **OK**.
The organization appears in the organization list and on the login page.
4. Click the organization to access the organization.

You are the superuser for the organization that you create. After you create an organization, you can configure the organization settings and invite users to access the organization. For more information, see ["Manage Organization Settings" on page 11](#) and ["Invite Users" on page 42](#), respectively.

Delete an Organization

You can delete an organization that you no longer manage or if you want to decommission the organization. You must be a user with the superuser role to delete an organization.



CAUTION: You cannot restore an organization after you delete it.

To delete an organization:

1. Log in to Juniper Data Center Assurance and navigate to **Organization > Settings**.
2. Click **Delete Organization**.
The Delete Organization page appears.
3. As a confirmation for deleting the organization, enter the name of the organization in the **Organization Name** field.
4. Click **Delete Organization**.
The organization is deleted and the Login page appears.

RELATED DOCUMENTATION

[Organization and Sites Overview](#) | 9

Manage Organization Settings

A superuser can configure the organization settings and do the following tasks:

- View organization name and organization ID, modify the organization name, and assign organization to a Managed Service Provider (MSP).
- Enable or disable the password policy for the organization and modify the password policy when the password policy is enabled.
- Modify the session timeout policy for the organization.
- Add, modify, and delete identity providers.
- Add, modify, and delete custom roles.
- Enable or disable the Juniper Networks support team access to the organization for troubleshooting.
- Configure webhooks for the organization.
- Add Juniper account to link Juniper Networks devices to the organization.
- Generate, edit, or delete API tokens used for authenticating users.

To configure and to manage organization settings:

1. Click **Organization > Settings** in the navigation menu.

The Organization Settings page appears.

2. Configure or modify the organization settings as needed. Refer to [Table 1 on page 11](#).
3. Click **Save** to save the settings.

Verify that the settings are saved and close the Organization Settings page.

Table 1: Organization Settings Parameters

Field	Description
Organization Name	Name of the organization. You can edit the organization name here.
Organization ID	The ID for the organization. The value is auto-generated. This is a read-only field.

Table 1: Organization Settings Parameters *(Continued)*

Field	Description
Managed Service Provider	Assign the organization to a Managed Service Provider (MSP), if any.
Password Policy	Enable or disable (default) password policy. If you enable the password policy, configure the password policy parameters; see Table 2 on page 13 .
Session Policy	Configure the time, in minutes, after which the logged in session with the application should timeout; see Table 3 on page 13 .
Identity Providers	View identity providers configured in the organization. Add, edit, or delete the identity providers; see "Manage Identity Providers" on page 16 .
Roles	View roles configured for SSO. Add, edit, or delete the roles; see "Manage Roles" on page 20 .
Webhooks	Webhooks enable you to get notifications when the events that you have subscribed for occur. Click to enable or disable (default) webhooks. If you enable webhooks, you must select the type of events for which you want to receive notifications; see Table 4 on page 13 .
API Tokens	Generate and view API tokens to authenticate users when they retrieve data by using REST APIs; see "Manage API Tokens" on page 22 .
Support Access	Enable (default) or disable the Juniper Networks support team access to the organization for troubleshooting. It is recommended to disable this feature except during specific time frames when you are working with support to resolve an issue. In such scenarios, temporarily enable access, and disable it when the issue is resolved. When this feature is enabled, the support personnel can see information about all the devices in the organization.
Juniper Account Integration	Add your Juniper Networks account to link your Juniper Networks devices to the organization; see Table 5 on page 15 .
gNMI Telemetry	Enable or disable gNMI Telemetry subscription.

Table 2: Parameters to Configure Password Policy

Field	Description
Required minimum password length	Enter the minimum number of characters that should be present in the password of a user's account. Default is 8 characters. Range: 8 to 32
Require special characters	Click to enable (default) or disable the use of special characters in the password.
Require 2-Factor Authentication	Click to enable or disable (default) two-factor authentication for users accessing the organization. If you enable two-factor authentication, a code is sent to an authenticator app. A user must enter the code in addition to the password to access an organization.

Table 3: Parameters to Configure Session Policy

Field	Description
Session Timeout (minutes)	Enter the number of minutes after which the session should timeout. Default is 20160 minutes.
Inactivity Timeout (minutes)	Enter the number of minutes of inactivity after which the session should timeout. Default is 0, indicating that the session does not time out because of inactivity. Range: 0 to 480 minutes

Table 4: Parameters to Configure Webhooks

Field	Description
Status	Select to enable or disable webhooks. The values are: - Enabled: Webhooks is enabled and you can get notifications on third-party applications when events you have subscribed to occur in the organization. - Disabled: Webhooks is not enabled and you cannot get notifications on third-party applications when events occur in your organization.

Table 4: Parameters to Configure Webhooks (*Continued*)

Field	Description
Webhook Type	Select the format in which notifications are to be sent when a subscribed event occurs. The options are: • HTTP POST • Splunk
Name	Enter the name of the server or application to which notifications for subscribed events are to be sent.
URL	Enter the URL of the server or application where the notifications are to be sent when a subscribed event occurs. You must configure webhooks to send notifications to third-party applications, when events you have subscribed to are triggered on the managed devices. To receive webhook notifications in a format that is compatible with the third-party application, you need to configure an intermediary that can interact with the sending and receiving applications. The recommended intermediary platform is Make. For more information, see "Configure Webhooks to Receive Event Notifications in Third-Party Applications" on page 341.
Topics	Select the events for which you want to receive webhook notifications.
Advanced Settings	
Verify Certificate	Enable or disable verification of certificates.
Secret	Enter the secret to validate that the notifications received are from valid hosts.
Custom Headers	
Key	Enter a unique key that the webhook endpoint can use to authenticate the event notifications.
Value	Enter a unique value for the key.

Table 5: Parameters to Add Juniper Account

Field	Description
Email Address	The e-mail address associated with your Juniper Networks account.
Password	The password associated with your e-mail address.

Authentication Methods Overview

IN THIS SECTION

- [Benefits of Single Sign-On | 16](#)

Juniper Data Center Assurance (DC Assurance) provides different authentication methods to authenticate users.

You can use one of the following authentication methods to log in to the application GUI.

- DC Assurance account—Users can create a DC Assurance account to access the application GUI.
- Social Sign-In—All users can enable Google social media sign-in (or single sign-on) on their user account page.
- Single Sign-On (SSO)—Superuser can configure third-party Identity Providers (IdP) to authenticate users in the organization.

While users have the necessary permission to configure and use their accounts to log in, administrators can configure Single Sign-On for users in the organization.

Super Users can create and manage users in an organization. User management includes inviting users to join an organization and revoking users' access to the organization.

You can use Google as an authentication provider to sign in to the application. Google sign-in uses OpenID Connect (OIDC) to authenticate users by verifying their Google account credentials. As an alternative, superusers can configure IdP in the Organization Settings page and map the default roles to the IdP profiles. Secure Assertion Markup Language (SAML 2.0) is supported for SSO authentication

using third-party IdPs. The IdP asserts a user's identity and allows the user to access the Web GUI based on the user's role.

Benefits of Single Sign-On

- Users can use a single account to log in to multiple platforms and applications.
- SSO simplifies password management for users and administrators through centralized authentication by IdP.

Manage Identity Providers

IN THIS SECTION

- [Add an Identity Provider | 16](#)
- [Edit an Identity Provider | 18](#)
- [Delete an Identity Provider | 18](#)
- [Field Descriptions | 18](#)

You (superuser) can configure third-party Identity Providers (IdPs) to enable Single Sign-On (SSO) across your Juniper Cloud application portals. You can use any IdP that supports Security Assertion Markup Language (SAML) 2.0.

In addition to configuring the IdP, you can set a preferred portal for your organization. The preferred portal setting determines which application portal users are redirected to after successful SAML-based authentication. By default, the portal in which the organization was created becomes the preferred portal. For more information, see "[Set a Preferred Portal for an Organization](#)" on page 27.

Add an Identity Provider

You must be a superuser to configure IdP in an organization.

You can add multiple IdPs to your organization from the Settings page. If you add multiple IdPs, ensure that a user account is associated with only one IdP.

To add an identity provider for SSO:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the **Add IDP** icon above the Identity Providers table.
The Add Identity Provider window appears.
3. Provide a unique name for the Identity Provider and click **Add**.
The Create Identity Provider page appears.
4. Configure the identity provider by using the guidelines in [Table 6 on page 19](#).



NOTE: Use the information from your IdP's SAML 2.0 configuration settings to fill the following fields:

- Issuer
- Signing Algorithm
- SSO URL

5. Click **Save** and close the window.
The identity provider is created and listed in the Identity Providers table.



REMEMBER: To ensure user access, you must:

- Create custom roles corresponding to the roles created in the IdP portal for the user. The custom roles take precedence over the roles assigned from the Administrators page. For more information, see ["Manage Roles" on page 20](#).
- Inform first time users to log in to the organization by using the SSO URL or an IdP-initiated login method. This is necessary to establish a user's account as an SSO account.

For subsequent logins, users can use either the SSO URL or the portal URL.

- Avoid using the same e-mail address for both SSO and local account.
- Delete any previously created local accounts once SSO is setup.

However, it is advised to retain at least one local user account with the superuser privileges. This ensures that if there is an issue with the SSO, such as an expired certificate, at least one administrator will have access to the organization.

Edit an Identity Provider

To edit an identity provider:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the identity provider you want to edit in the Identity Providers table.
The Edit Identity Provider page appears.
3. Edit the identity provider by using the guidelines in [Table 6 on page 19](#).



NOTE: You cannot edit identity provider type, ACS URL, and Single Logout URL.

4. Click **Save**.
You are returned to the Organization Settings page, where you can view the changes in Identity Providers table.

Delete an Identity Provider

After you delete an identity provider, a user can log in only by using their local Juniper Data Center Assurance account.

To delete an identity provider:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the identity provider that you want to delete.
The Edit Identity Provider page appears.
3. Click **Delete**.
You are returned to the Organization Settings page, where you can view that the identity provider is removed from the Identity Provider table.

Field Descriptions

[Table 6 on page 19](#) lists the parameters to add identity providers to an organization.

Table 6: Parameters to Add Identity Providers

Field	Description
Name	Enter a name for the identity provider.
Type	Displays the type of identity provider. The default identity provider is SAML and cannot be modified.
Issuer	Enter the unique URL that identifies your SAML identity provider. You can get this information from your IdP's SAML 2.0 configuration settings.
Name ID Format	Select the unique identifier for the user. The options are e-mail and unspecified. E-mail—The identity provider uses your e-mail address to authenticate you. Unspecified—The identity provider generates a unique identifier to authenticate you.
Signing Algorithm	Select the algorithm used to sign SAML assertions • SHA1 • SHA256 (default) • SHA384 • SHA512 You can get this information from your IdP's SAML 2.0 configuration settings.
Certificate	Enter the certificate issued by the SAML identity provider. NOTE: Download the certificate from the IdP dashboard, copy the entire text of the certificate, including the header and footer, and paste it in the field.
SSO URL	Enter the URL to redirect the users to the SAML identity provider for authentication. You can get this information from your IdP's SAML 2.0 configuration settings.
Custom Logout URL	Enter the URL to redirect the users after logging out.

Table 6: Parameters to Add Identity Providers (*Continued*)

Field	Description
ACS URL	The Assertion Consumer Service (ACS) URL where the IdP sends the authentication response after the user signs in. The value is auto-generated and not editable. Copy the ACS URL and store it securely. You will need the URL to complete the SAML 2.0 integration in your IdP portal.
Single Logout URL	The URL that the identity provider should redirect when a user logs out of an authentication session. The value is auto-generated and not editable. Copy the Single Logout URL and store it securely. You will need the URL to complete the SAML 2.0 integration in your IdP portal.

Manage Roles

IN THIS SECTION

- [Add a User-Defined Role | 21](#)
- [Edit a User-Defined Role | 21](#)
- [Delete a User-Defined Role | 22](#)

A user with the Super User role can create a new role that maps a user role in an enterprise to a pre-defined role in Juniper Data Center Assurance.

For example, you can configure an administrator role and map it to the Network Admin role so that the administrator role has the access privileges of the Network Admin user. The Network Admin role can be assigned to any enterprise user. [Table 7 on page 21](#) lists the parameters to add custom roles to an organization.

Table 7: Parameters to Add Roles

Field	Description
Name	Enter a name for the role.
Role	Select an access level for the role: • Super User • Network Admin • Observer (default) • Helpdesk NOTE: This role is not currently supported. See "Predefined User Roles Overview" on page 38 for details on privileges of each role.

Add a User-Defined Role

A Super User can add a user-defined role and map it to a pre-defined role in Juniper Data Center Assurance.

To add a user-defined role that maps to a pre-defined role:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the **Create Role** icon.
The Create Role page appears.
3. Configure the new role by following the guidelines in [Table 7 on page 21](#).
4. Click **Create**.
The new role is listed in the Roles table.

Edit a User-Defined Role

To edit a user-defined role:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the role that you want to edit.
The Edit Role page appears.
3. Edit the name and role by following the guidelines in [Table 7 on page 21](#).
4. Click **Save**.
You are returned to the Organization Settings page, where you can verify the changes in the Roles table.

Delete a User-Defined Role

After you delete a user-defined role, users assigned to the user-defined role must be assigned one of the roles defined in Juniper Data Center Assurance to continue accessing the resources in Juniper Data Center Assurance.

To delete a user-defined role:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the role that you want to delete.
The Edit Role page appears.
3. Click **Delete**.
You are returned to the Organization Settings page, where you can verify that the custom role is not listed in the Roles table.

Manage API Tokens

IN THIS SECTION

- [Add an API Token | 23](#)
- [Edit an API Token | 24](#)
- [Delete an API Token | 24](#)

API tokens use REST APIs to authenticate users when they try to retrieve information from Juniper Data Center Assurance. By using API tokens, users can avoid authentication for each request they make. An API token provides visibility into the resources accessed by a user, enabling you to have better control over access to resources.

[Table 8 on page 23](#) lists the parameters for configuring API tokens.

Table 8: Parameters to Configure API Tokens

Field	Description
Name	Name of the API token.
Role	Role to which the API token is applicable: • Super User • Network Admin • Observer • Helpdesk NOTE: This role is not currently supported.
Key	The key auto-generated to identify the application the user is using to access the resources.

Add an API Token

To add an API token for a role:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the **Create Token** icon.
The Create API Tokens page appears.
3. Enter values by following the guidelines in [Table 8 on page 23](#).
4. Click **Generate**.
The API token is populated in the **Key** field.
5. Click **Close** to return to the Organization Settings page.

Edit an API Token

To edit an API token:

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the API token that you want to edit.
The Edit API Token page appears.
3. Edit the name, role, and site access by following the guidelines in [Table 8 on page 23](#).
4. Click **Save**.
You are returned to the Organization Settings page, where you can verify the changes in the API Tokens table.

Delete an API Token

To delete an API token:



NOTE: Users using API tokens to access Juniper Data Center Assurance resources cannot access the resources after the API token is deleted.

1. Click **Organization > Settings** in the navigation menu.
The Organization Settings page appears.
2. Click the API token that you want to delete.
The Edit API token page appears.
3. Click **Delete**.
You are returned to the Organization Settings page, where you can verify that the API token is not listed in the API Tokens table.

Integrate Your Juniper Support Resources to Your Organization

To enable the correlation of device information maintained within Juniper's support databases with your Juniper Support Insights experience, you must associate your Juniper support resources with your organization. Use your Juniper Support credentials (created through the [Juniper Support Portal](#)) to link the Juniper account and integrate your support resources to your organization. Once the account is

successfully linked to the organization, the Installed Base (**Organization > Inventory > Installed Base**) page is populated with support resources for the devices associated with the account.

When you link your Juniper-registered account to an organization, all primary and secondary assets associated with the account and covered by an active service contract are integrated with the organization.



NOTE: If the Juniper account is linked to multiple organizations, a device onboarded to one organization is automatically removed from the Installed Base page of the other organizations. However, once the device is offboarded from the organization, the device reappears on the Installed Base page of all the organizations.

For more information on device specific details collected from the cloud-connected devices, see ["About the Inventory Page" on page 51](#).

To integrate your Juniper support resources to your organization:

1. Click **Organization > Settings** to open the Organization Settings page.




NOTE: If no Juniper account is currently associated with the organization, the Installed Base tab on the Inventory page will display a link to add a Juniper account. Click **Add Juniper Account** to link the Juniper-maintained asset information you are authorized to access with your organization.

You can associate multiple Juniper accounts with your organization. This allows you to view a consolidated list of devices associated with multiple accounts in one place.

The Add Juniper Account pop-up also displays the list of primary accounts that are currently linked to the organization.

Locate the **Juniper Account Integration** tile.



The image shows a web interface titled "Juniper Account Integration". Below the title is the text "Link your Juniper account to access your Juniper Networks assets". There is a button labeled "Add Juniper Account". Below the button is a table with two columns: "Name" and "Actions". The table is currently empty.

2. On the Juniper Account Integration tile, click **Add Juniper Account**.

The Add Juniper Account window appears.

3. Enter your Juniper Support credentials (e-mail address and password), and then click **Add**.

Juniper Data Center Assurance validates your credentials and links the Juniper account with the organization.

The Juniper Account Integration tile displays the account that is linked to the organization.



NOTE: Once you successfully execute the account integration workflow, your identity is automatically validated and you can access the device-centric and customer-centric JSI data artifacts of devices associated with the linked Juniper accounts, both the primary account and all the secondary accounts. Otherwise, you will be prompted to validate your identity to gain access to the customer-centric JSI data artifacts of the devices. Once validated, validation persists across user sessions, eliminating the need to re-validate your identity each time you log in. For more information, see "[Validate Your Juniper Account](#)" on page 69.

The Installed Base (**Organization > Inventory > Installed Base**) page is populated with the support resources for devices associated with the account.



NOTE: To remove an account, click the Delete (trash can) icon against the account name on the **Juniper Account Integration** tile. When you remove an account, the associated devices are removed from the **Installed Base** page.

Set a Preferred Portal for an Organization

Juniper offers multiple cloud application portals to cater to your various business needs, such as [Juniper Routing Assurance](#), [Juniper Data Center Assurance](#), and [Juniper Support Insights \(JSI\)](#). In a multi-portal environment, the preferred portal setting allows you to specify the primary application portal users are directed to when accessing an organization.

When users log in to an organization through a Global Cloud Portal or an Identity Provider (IdP), the preferred portal setting determines which application portal the users are redirected to after authentication.

When an organization is created, the portal used to create the organization becomes the preferred portal. During logins, users are redirected to this portal after successful authentication.

Superusers can update the preferred portal at any time from the Organization Settings (**Organization > Settings**) page. Once updated, all future Global Cloud Portal or IdP logins for that organization will redirect users to the newly selected portal.

For example, if you create an organization from the Juniper Support Insights (JSI) portal, users logging in through the IdP will be redirected to the JSI portal, by default. If you later subscribe to Juniper Routing Assurance (JRA) and wants to use it as the primary portal for the organization, you can update the Preferred Portal setting to the JRA portal. This ensures that future logins for the organization are redirected to the JRA portal instead of JSI.



NOTE: This setting is shared across applications but applies only to the organization in which it is configured.

To update the preferred portal for an organization:

1. Click **Organization > Settings** in the navigation menu.

The Organization Settings page appears.

2. Locate the **Preferred Portal** drop-down.

By default, the application in which the organization was created appears pre-selected.

3. Select the desired application portal from the drop-down to update the preferred portal.

Your options are:

- Routing Assurance
- JSI
- Apstra

4. Click **Save** to apply the change.

The selected portal is now set as the preferred portal for this organization.



NOTE: The preferred portal setting applies only to Global Cloud Portal and IdP-based logins and does not apply to local portal logins.

All future logins for this organization will redirect users to the selected portal.

3

CHAPTER

Site Management

IN THIS CHAPTER

- [About the Sites Page | 30](#)
 - [Edit Site Configuration | 31](#)
-

About the Sites Page

IN THIS SECTION

- [Tasks You Can Perform | 30](#)
- [Field Description | 30](#)

Sites are the physical locations that host devices, such as data centers within an organization's network. In Juniper Data Center Assurance, a site is representation of a blueprint deployed in Apstra. Whenever a new blueprint is deployed in Apstra, a new site corresponding to the blueprint is created automatically in Juniper Data Center Assurance.

To access the Sites page, click **Organization > Site Configuration**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about the sites in an organization—You can view the site name, country, time zone, address, the site group the site belongs to, and notes about the site.
- Modify site configuration; ["Edit Site Configuration" on page 31](#).
- Filter the data displayed in the table—You can filter site information using keywords. Enter the search term in the text box next to the search icon (magnifying glass) and press **Enter**. The search results are displayed on the same page. You can also add one or more filters and clear filters.

Field Description

[Table 9 on page 30](#) describes the fields displayed on the Sites page.

Table 9: Fields on the Sites Page

Fields	Description
Name	Displays the name of the site.
Country	Displays the country where the site is located.

Table 9: Fields on the Sites Page (*Continued*)

Fields	Description
Timezone	Displays the time zone of the site.
Address	Displays the address of the site.
Site Groups	Displays the site groups to which the site belongs, if any.
Notes	Displays additional information about the site.

Edit Site Configuration

In Juniper Data Center Assurance, a site represents a blueprint deployed in Apstra Data Center Director. When a blueprint is deployed, a site corresponding to the blueprint is created automatically in Juniper Data Center Assurance. After a site is created, superusers can update its settings as required from Juniper Data Center Assurance.

Superusers can modify the configuration of an existing site from the Site Configuration: *site-name* page.

To modify site configuration:

1. Click **Organization > Site Configuration**.

The Sites page appears.

2. Click the site you want to edit.

The Site Configuration: *site-name* page appears.

3. Modify the required site configuration, see [Table 10 on page 31](#).

4. Click **Save**.

A confirmation message indicating that the site configuration is updated is displayed.

Table 10: Fields on the Site Configuration: *site-name* Page

Fields	Description
Site Name	Name of the site. The site name can contain up to 64 characters.

Table 10: Fields on the Site Configuration: *site-name* Page (Continued)

Fields	Description
Location	Select the site's location by clicking on the map, entering the coordinates (latitude and longitude), or searching for the location in the search field. Selecting a location automatically updates the fields for country and time zone.
Country	Select the country where the site is located. If the location is selected on the map, this field is automatically updated with the respective country. However, if you select a country from the drop-down list, the same is not automatically reflected on the map.
Time Zone	Select the timezone of the site. If the location is selected on the map, this field is automatically updated with the respective timezone. However, if you select a country from the drop-down list, the same is not reflected on the map.
Notes	Enter additional information about the site. Notes can contain up to 1000 characters.
Site Groups	Select the site groups to which the site should belong, if any. If no site group is available, you can type a name for the site group and press Enter to create the site group.

Table 10: Fields on the Site Configuration: *site-name* Page (Continued)

Fields	Description
Webhooks	Webhooks enable you to get notifications on third-party applications when the events that you have subscribed for occur. If you enable webhooks, you must select the type of events for which you want to receive notifications. Webhooks configured at the site level apply only to events generated for that site. Fields for configuring webhooks: • Status: Select to enable or disable webhooks. The values are: • Enabled: Webhooks is enabled and you can get notifications on third-party applications when events you have subscribed to occur in the organization. • Disabled: Webhooks is not enabled and you cannot get notifications on third-party applications when events occur in your organization. • Webhook Type: The format in which the notifications are sent to third-party applications. • HTTP POST • OAuth 2.0 • Splunk • URL: Enter the URL of the server or application where the notifications are to be sent when a subscribed event occurs. • Topics: Select the events for which you want to receive webhook notifications. Advanced Settings • Verify Certificates: Enable or disable certificate verification. • Secret: Enter the secret to validate that the notifications received are from valid hosts. • Custom Headers: Webhook custom headers are key-value pairs that provide additional information about the notifications. You can add multiple custom headers to:

Table 10: Fields on the Site Configuration: *site-name* Page (Continued)

Fields	Description
	• Provide additional information in plain text, along with the default headers, about the webhook notifications being sent to the configured endpoint. • Provide security, such as API keys, to verify end-to-end data integrity, for authorization, and so on. Click the Add Header to add webhook headers. • Key—Enter a unique key. • Value—Enter a unique value for the key. The value can contain alphanumeric characters. Click the Delete icon (trash can) to remove the webhook headers.
Sustainability	Configure carbon intensity and average power cost for a site to ensure that the Site-Level Energy Dashboard displays accurate GHG emissions and power cost. To configure the Average Power Cost, enter the cost you pay in USD for one kilowatt-hour of electricity as a numeric value between 0.00 and 1.00. To configure Carbon Intensity, enter the number of grams of CO₂e emitted per kilowatt-hour (gCO₂e/kWh) for the site. The value must be a non-negative numeric value. You can obtain this value from your power supplier, regional energy authority, or your organization's sustainability team. For more information on the sustainability tile, see "Configure Sustainability Settings for a Site" on page 227.

Table 10: Fields on the Site Configuration: *site-name* Page (Continued)

Fields	Description
Port Threshold	Adjust the slider to set the port utilization thresholds. The slider has two draggable handles. Left Handle—Drag the handle to control the boundary between Low and Medium utilization. Move this handle to the right to increase the Low utilization range (light green band expands). Move this handle to the left to increase the Medium utilization range (medium green band expands). Right Handle—Drag the handle to control the boundary between Medium and High utilization. Move this handle to the right to increase the Medium utilization range (medium green band expands). Move this handle to the left to increase the High utilization range (dark green band expands). By default, the threshold for low traffic is set at 10% and high traffic at 80%. The slider displays the percentage at each handle position, indicating how much of the scale each utilization range occupies. For more information, see "Configure Port Threshold for a Site" on page 282.

4

CHAPTER

User Management

IN THIS CHAPTER

- [About the Administrators Page | 37](#)
 - [Predefined User Roles Overview | 38](#)
 - [Add Users to an Organization | 40](#)
 - [Invite Users | 42](#)
 - [Manage Users and Invites | 44](#)
 - [Manage Your Account | 46](#)
-

About the Administrators Page

IN THIS SECTION

- [Tasks You Can Perform | 37](#)
- [Field Descriptions | 37](#)

To access the Administrators page, click **Organization> Administrators** in the navigation menu.

Tasks You Can Perform

An administrator with the superuser role can perform the following tasks from this page:

- View details of the existing users and the users who are invited to access the organization—The basic information about the users, such as first name, last name, e-mail address, invite status of the user, and role assigned is displayed. See [Table 11 on page 37](#) for field descriptions.
- Invite administrators; see ["Invite Users" on page 42](#).
- Manage administrator invitations; see ["Manage Users and Invites" on page 44](#).

Field Descriptions

[Table 11 on page 37](#) describes the fields on the Users page.

Table 11: Fields on the Users Page

Fields	Description
First Name	The first name of the user.
Email	The e-mail address the user would use to access Juniper Data Center Assurance.

Table 11: Fields on the Users Page *(Continued)*

Fields	Description
Status	Indicates a user's account status: • Active: The user's account is active and the user can access the organization. • Invite Pending: The user is yet to accept the e-mail invitation sent to them and doesn't have access to the organization or the user has rejected the invitation to access the organization. • Invite Expired: The e-mail invitation sent to the user has expired. An invitation expires after seven days.
Role	The role assigned to a user. See "Predefined User Roles Overview" on page 38 for details about the user roles.

RELATED DOCUMENTATION

| [Add Users to an Organization](#) | 40

Predefined User Roles Overview

IN THIS SECTION

- [Administrator Roles](#) | 39
- [Limited Roles](#) | 40

Juniper Data Center Assurance provides four predefined administrator roles and three predefined limited roles to manage access privileges of users, based on the tasks they need to perform.

A superuser creates an organization, adds users to predefined roles depending on the requirements of the organization. For example, an organization with a large number of networking devices would require

multiple users performing different roles to efficiently manage the organization, whereas, in a small organization, a single user can perform the tasks to be carried out by users with all the roles.

Different types of users in an organization, such as a network architect, network planner, NOC engineer, and field technician, all derive their access privileges from the predefined roles assigned to them.

The following sections outline the various user roles and their access privileges.

Administrator Roles

User Roles and their Responsibilities

The four predefined roles are:

- Super User
 - Is the administrator of the organization.
 - Manages organization-level settings and user access.
 - Creates organization, invites users, assigns user roles, creates sites, adopts devices, and so on.
 - Does not require high-level of networking domain expertise.
- Network Admin
 - Is a networking expert who monitors, verifies, and troubleshoots within the assigned scope of an organization's network.
 - Has full access to the sites and site groups assigned to their role.

If no sites or site groups are assigned, the Network Admin is granted organization-level access by default.

 - Can take corrective actions, perform configurations, and review performance metrics.
- Observer
 - Monitors network events across the sites or site groups assigned to their role.

If no sites or site groups are assigned, the Observer is granted organization-level read-only access by default.

 - Cannot take corrective action; escalates issues to the network administrator for resolution.
- Helpdesk
 - Monitors events for assigned sites and site groups in the organization.



NOTE: This role is not currently supported.

Limited Roles

User Roles and their Responsibilities



NOTE: The access privileges are limited for these role.

The predefined limited roles are:

- Reporting
 - Can access all the analytics tools.
- DC Edge Admin
 - Can onboard data center edges and monitor device status during onboarding for a defined grace period set by the superuser.
- Super Observer
 - Monitor all sites within the organization.
 - Has read-only access to organization settings.


Add Users to an Organization

An administrator with the Super User role can add users to an organization and provide role-based access by sending an invitation to the user's e-mail ID. The user needs to accept the invitation to be a member of the organization.

Existing users can access their organization by using their Juniper Data Center Assurance account.

[Figure 2 on page 41](#) illustrates the workflow for inviting a new user to an organization.

Figure 2: Add users to an organization



The status of the invitation is shown as Invite Pending until the user:

- Accepts the invitation to get role-based access to the organization.
- Rejects the invitation to access the organization.
- Doesn't accept or reject the invitation within seven days. The status of such invitations is displayed as Invite Expired.

If the user accepts the invitation and has role-based access to the organization, but if the superuser wants to take away the user's access, the superuser can revoke the invitation.

If the user invitation expires, the superuser can re-invite the user or cancel the invitation.

Invite Users

An administrator with the Super User role can add users to an organization by sending an e-mail invitation from the Administrators page.

The user must accept the invitation within seven days, after which the invitation expires.

A user's access privileges within the organization is based on the role assigned to the user. A user can be assigned only one role in an organization in Juniper Data Center Assurance. However, a user can be a member of multiple organizations within Juniper Data Center Assurance, and can have different roles in each organization. For more information on roles, see ["Predefined User Roles Overview" on page 38](#).

To invite a user:

1. Click **Organization > Administrators**.

The Administrators page appears.

2. Click the **Invite Administrators** icon.

The Administrators: New Invitation page appears.

3. Enter user details and assign a role according to the guidelines provided in [Table 12 on page 42](#).

4. Click **Invite**.

A confirmation message indicating that the user is invited is displayed, and the user details are listed on the Users page.

5. Check the status of the user. The status is displayed as Invite Pending. If the status changes to Invite Expired, you can delete the user, reinvite the user or cancel the invitation. For more information, see ["Cancel an Invitation" on page 45](#) and ["Reinvite a User" on page 45](#).

Table 12: Fields on the Invite User Page

Field	Description
First Name	Enter the first name of the user. First name can contain up to 64 characters.
Last Name	Enter the last name of the user. Last name can contain up to 64 characters.
Email	The e-mail address the user would use to access Juniper Data Center Assurance.

Table 12: Fields on the Invite User Page *(Continued)*

Field	Description
Role	Assign a role to the user. You can assign only one role to a user in an organization. You can assign: Administrator Roles • Super User • Network Admin • Observer • Helpdesk NOTE: This role is not currently supported. Limited Roles • Reporting • DC Edge Admin • Super Observer See "Predefined User Roles Overview" on page 38 for information about user roles.
Site Access	Select the sites the user can access in the organization. You must select at least one site. The options are: • All Sites: The user can access all sites in the organization. • Site Groups: The user can access specific site groups. Click Add (+) icon to add one or more site groups from the drop-down list. • Specific Sites: The user can access specific sites in the organization. Click Add (+) icon to add specific sites from the drop-down list.

Manage Users and Invites

IN THIS SECTION

- [Edit User Role | 44](#)
- [Reinvite a User | 45](#)
- [Cancel an Invitation | 45](#)
- [Revoke a User | 46](#)

You must be an administrator with the Super User role to manage users and user invitations. You can edit user role, reinvite, cancel invitations, and revoke users from the Administrators page.

Edit User Role

On the Administrators: *Name* page, you can edit the role of a user. The first name, last name, and e-mail ID of a user cannot be modified.

To edit user role:

1. Click **Organizations > Administrators**.

The Administrators page appears.

2. Select the user whose role or site access you want to edit.

The Administrators: *Name* page appears.

3. Modify the role or site access as needed. See ["About the Administrators Page" on page 37](#) for more information about user roles.



NOTE:

- If you modify the role or site access of a user whose invitation status is Active, the user is not notified about the modification in the role.

- If you modify the role or site access of a user whose invitation status is Invite Pending or Invite Expired, a new invitation e-mail is sent to the user to access the organization with the new role-based access privileges.

4. Click **Save**.

A confirmation message indicating that the user invitation is updated is displayed and you are returned to the Administrators page, where you can view the changes you made.

Reinvite a User

You can reinvite a user if:

- The user invitation expired.
- The user invitation is pending.
- The user role or site access needs to be modified for users with Invite Pending or Invite Expired invitation status.

To reinvite a user to the organization:

1. Click **Organizations > Administrators**.

The Administrators page appears.

2. Select the user you want to reinvite and click **Re-invite**.

You can reinvite a user whose status is Invite Expired or Invite Pending. For users whose access is revoked or deleted, you must click the **Invite User (+)** icon to reinvite the user; see ["Invite Users" on page 42](#).

An invitation e-mail is sent to the user and the user account is listed on the Administrators page with status Invite Pending.

If the user doesn't accept the invitation within seven days, the invitation expires.

Cancel an Invitation

You can invalidate an invitation by canceling the invitation. You can uninvite an administrator if the invitation status is Invite Pending or Invite Expired on the Administrators page.



NOTE: An invite expires after seven days.

To uninvite a user:

1. Click **Organizations > Administrators**.

The Administrators page appears.

2. Select the user you want to uninvite and click **Uninvite**.

A confirmation message indicating that the invite is canceled is displayed and you are returned to the Administrators page. The details about the user invitation is no longer listed in the Administrators table.

Revoke a User

If the user accepts the invitation and has role-based access to the organization, but you want to take away the user's access, you can revoke the invitation. Revoking a user's access deletes the user from the organization. You can revoke access only for active accounts.

To revoke a user's access to an organization:

1. Click **Organizations > Administrators**.

The Administrators page appears.

2. Select the user whose access needs to be revoked and click **Revoke Access**.

The user is deleted from the organization and cannot access the organization.



NOTE: Juniper Data Center Assurance maintains a log of the user's activities in the organization even after the user's account is deleted or their access gets revoked. For example, the user's activities recorded in the audit logs will remain even if they no longer have access to the organization.

Manage Your Account

You can manage your account information from the My Account page. You can access the My Account page by clicking the user account icon in the top right corner of the GUI. From the list, choose **My Account**.

You can perform the following tasks in the My Accounts page:

- [Change account information](#)
- [Change your password](#)

- [Enable two-factor authentication](#)
- [Set time preference and MAC address format](#)
- [Enable e-mail notifications for superusers and network administrators](#)
- [Enable social sign-in](#)
- [Create an organization](#)
- [Delete your account](#)

1. To change account information:

- a. Click your user account icon at the top-right corner and click **My Account** from the list.
- b. Change your e-mail address, name, and phone number, as necessary, in the Account Information section.
 To change your e-mail address, click **CHANGE**. In the Change Email window, enter the new e-mail address.
 A verification link is sent to the new e-mail address and the address is updated in the Juniper Data Center Assurance after authentication.
- c. Click **Save**.
 Your user account information is successfully updated.

2. To change your password:

- a. Type a password in the New Password box.
 The superuser configures the password policy for the organization. A password can contain up to 32 characters including special characters.
- b. Click **Save**.
 A message confirms that your user data is successfully updated.

3. To enable two-factor authentication:

- a. Enable **Two Factor Authentication** check box, under Authentication.
- b. Click **Save**.
 A message confirms updating your user data. A verify button appears near the two-factor authentication option.
- c. Click **Verify**.
 The Verification of Two Factor Authentication page displays a QR code.
- d. Open your authenticator application and click the add icon (+) to add a new account.
- e. Scan the QR code displayed.

Your account appears in your authenticator application.

- f. Enter the token number from your authenticator application in the Verification of Two Factor Authentication page.

- g. Click **Verify**.

A green check mark appears beside the Two Factor Authentication option on your My Account page. The two-factor authentication is active for your account. You can log out and log back in to the cloud portal.

4. To set your time preference and MAC address format:

You can format the way time and MAC address is displayed on the GUI.

- a. Under Preferences, select the MAC address format you prefer. The three options are xx:xx:xx:xx:xx:xx, xx-xx-xx-xx-xx-xx, xxxx.xxxx.xxxx.
- b. Under Preferences, select 24-Hour Time if you prefer that time format, or leave the box unchecked if you prefer a 12-hour time format.

- c. Click **Save**.

5. To enable e-mail notifications:

You must enable e-mail notification on the My Account page to receive e-mail notifications for all or selected sites. You can also enable e-mail notifications at the organization level.

- a. Click **Enable** in the Email Notification section.

The Enable Email Notifications page appears.

- b. Click the **Enable Org Notifications** toggle button to enable e-mail notifications at the organization level.

- a. Click the toggle button against a site to receive e-mail notifications specific to the site.

- b. Click **Close**.

The Enable Email Notification section shows that you have enabled notifications for your current organization.

6. To enable social sign-in:

- a. Enable the Sign In With Google option in the Social Sign In section.

A message asks your permission for redirection to link your Google account.

- b. Click **Yes**.

You will be redirected to the Google sign in page.

- c. Enter your Google e-mail and password and click **Next**.

Juniper Data Center Assurance links your Google account and redirects to the My Account page. A message confirms that Juniper Data Center Assurance has linked your Google account.

7. To create an organization:

- a. Click **Utilities > Create Organization**.

The Create Organization window appears.

- b. Enter a name for the organization.

- c. Click **OK**. The new organization opens in a new tab.

8. To delete your account:

- a. Click **Utilities > Delete Account**.

A confirmation message appears.

- b. Click **Yes**.

9. Link your local account with global account:

- a. On the **Link to Global Account** tile, click **LINK**.

The Link to Global Account window appears.

- b. Click **Link**.

For more information on linking accounts, see No Link Title.

5

CHAPTER

Inventory Management

IN THIS CHAPTER

- [About the Inventory Page | 51](#)
 - [DC Edges and DC Switches | 64](#)
 - [Validate Your Juniper Account Overview | 67](#)
 - [Validate Your Juniper Account | 69](#)
 - [Configure Ports | 75](#)
-

About the Inventory Page

IN THIS SECTION

- [Tasks You Can Perform | 51](#)
- [Field Description for Installed Base Page | 58](#)

The Inventory page lists the devices in an organization.

To access the Inventory page, click **Organization** > **Inventory** on the navigation menu.

Tasks You Can Perform

You can perform the following tasks on the Inventory page:

- View Juniper-maintained support information for all the Juniper Networks devices associated with your organization, whether they are directly connected to the cloud or linked to your organization through the account integration workflow from the **Installed Base (Organization > Inventory > Installed Base)** page.



NOTE: To access Juniper-maintained asset information about the Juniper Networks devices from the **Installed Base** page, you (superuser) must first link your Juniper account to your organization.

To integrate your Juniper account with the organization:

1. Click **Add Juniper Account** on the Installed Base page.

The Add Juniper Account pop-up appears.

2. Enter your Juniper support credentials (created through the [Juniper Support Portal](#)).
3. Click **Add**.

On successful integration, the primary account associated with your credentials is listed in the Linked Accounts section. The Linked Account section also lists other primary accounts that are currently linked to the organization.

Alternatively, you can link your account to the organization from the Settings page (**Organization > Settings > Juniper Account Integration**). If you no longer wish to view Juniper-maintained asset information for devices associated with an account in your organization, you can unlink the account from this page. This removes the devices associated with the account from the organization. For more information, see ["Integrate Your Juniper Support Resources to Your Organization" on page 24.](#)

Use the Installed Base tab to:

- View all Juniper assets associated with your organization and track their operational status
- Identify devices nearing hardware Hardware End of Life, Sale, and Support milestones
- Track software FRS, End of Engineering, Support milestones
- Identify devices running outdated software and upgrade to recommended Junos OS release
- Open the Device Details page to access a device's security advisories, bug notifications, hardware and software milestones, contract details, and so on.

Click a row in the device table to navigate to the JSI Device Details page.

- Export the details as a CSV file for offline analysis or customization.

The Installed Base includes an insights panel displaying the following data:




Table 13: Insights Panel Field Description

Category	Description	Usage Notes
Device Models	Number of devices per model.	Use it to identify the distribution of device models in your organization.
Onboarding	The percentage of devices in each connection state. - Assured—Devices that are onboarded to the Assured service level. This service level includes access to the JSI experience. - Onboarded—Devices that are currently onboarded to the JSI service-level only. - Not Onboarded—Devices that are not onboarded to any service-level.	Track the connection status of devices in the organization.
Hardware – End of Support (EOS)	Percentage of devices nearing or past hardware EOS. - Expired—Devices that have reached their EOS. 0 - 3 months—Devices whose EOS is in less than 3 months. 0 - 12 months—Devices whose EOS is in less than 12 months. > 12 months—Devices whose EOS is after 12 months.	Identify devices nearing or past hardware EOS to plan renewals or replacements.

Table 13: Insights Panel Field Description (*Continued*)

Category	Description	Usage Notes
Hardware – End of Life (EOL)	Percentage of devices nearing or past hardware EOL. - Expired—Devices that have reached their EOL. 0 - 3 months—Devices whose EOL is in less than 3 months. 0 - 12 months—Devices whose EOL is in less than 12 months. > 12 months—Devices whose EOL is after 12 months.	Identify devices nearing EOL to plan hardware migrations.
Contract State	The percentage of Supported and Unsupported devices. Supported—Devices recognized as being covered by an active support contract. These devices are eligible for JTAC support. Unsupported—Devices recognized as not being covered by an active support contract. JTAC support may not be available for these devices until the contract is renewed. NOTE: Unsupported devices can still be onboarded and managed within your organization. For more information on device contracts, see "About the Contracts Page" on page 80.	Monitor support eligibility, prioritize renewals, and filter the Device table by state. Supported filter is selected by default. You can click it to clear the selection.

Table 13: Insights Panel Field Description (*Continued*)

Category	Description	Usage Notes
Software – End of Support (EOS)	Percentage of devices running software nearing or past EOS. - Expired—Devices running software that have reached EOS. 0 - 3 months—Devices running software whose EOS is in less than 3 months. 0 - 12 months—Devices running software whose EOS is in less than 12 months. > 12 months—Devices running software whose EOS is after 12 months.	Identify devices running software past or nearing software EOS and plan upgrades to recommended releases.
Software – End of Engineering (EOE)	Percentage of devices running software nearing or past EOE. - Expired—Devices running software that have reached EOE. 0 - 3 months—Devices running software whose EOE is in less than 3 months. 0 - 12 months—Devices running software whose EOE is in less than 12 months. > 12 months—Devices running software whose EOE is after 12 months.	Identify devices approaching EOE to plan upgrades to recommended releases.

For information on the fields on the Installed Base device table, see ["Field Description for Installed Base Page" on page 58](#).

- **Access support insights Device Details Page**

You can view detailed support insights for the device in the Device Details page.

To access the Device Details page for a device, click the device entry in the table on the Installed Base tab. The Device Details page appears.



NOTE: JSI masks data to prevent unauthorized user access to customer-centric support insights associated with the device. You must successfully validate your identity to access the masked customer-centric information.

To validate your identity:

1. Do one of the following:

- Hover over the blurred fields and click **Validate**.
- Click **Validate Email ID** on the banner.

The Validate Juniper Account pop-up appears.

2. Enter your Juniper registered account credentials (email address and password) that you created through [Juniper Support Portal](#) and click **Validate**.

On successful validation, the page is refreshed and displays the customer-centric information that you are entitled to access.

For more information on validating your identity, see "[Validate Your Juniper Account](#)" on page 69.

The Device Details page displays the following information:

- **DEVICE VULNERABILITIES**

- **Security Vulnerabilities**—Click the link to view information on the security vulnerabilities advisories published by the Juniper Security Incident Response (SIRT) team for the device.



NOTE: If the Juniper Networks device is in cloud-connected (Assured or Onboarded) state, the **Security Incident Response Team Tickets** page displays a list of security vulnerabilities specific to the type of Juniper Networks device and the Junos OS version installed.

But if the Juniper device is not connected to cloud, the **Security Incident Response Team Tickets** page displays a pop-up indicating that you must connect the device to cloud to view the list of security vulnerabilities specific to the device.

Click **OK** to be redirected to the Installed Base page.

The security vulnerabilities page for a device also displays a banner with the total counts and percentages of critical, high, medium, and low severity vulnerabilities for the device. You can filter the advisories relevant to the device based on the severity level. Click a severity level criterion on the banner to filter the advisories for that level. For example, if you click the

Critical count or the Critical Vulnerabilities percentage criteria, the advisories are filtered to display only the critical advisories for the device.

From the Security Incident Response Team Tickets page, select an entry to view the **Device SIRT Quick View** pane. The Device SIRT Quick View pane displays detailed information about the advisory.

For more information on the fields on the Security Incident Response Team Tickets page and Device SIRT Quick View pane, see [Table 15 on page 60](#).

- **Pro-Active Bug Notifications**—Click the link to view information on the known software issues impacting the device.



NOTE: If the Juniper Networks device is cloud-connected (Assured or Onboarded), the **Proactive Bug Notifications** page displays a list of known software issues relevant to the type of Juniper Networks device and the Junos OS version installed.

But if a Juniper device is not connected to cloud, the **Proactive Bug Notifications** page displays a pop-up indicating that you must connect the device to cloud to view the list of known software issues specific to the device.

Click **OK** to be redirected to the Installed Base page.

The Proactive Bug Notifications page displays a banner with the total counts and percentages of critical, major, and minor known software issues for the device. You can filter Proactive Bug Notifications for a device based on the risk level. Click a risk level criterion on the banner to filter the Proactive Bug Notifications for that level. For example, if you click the Critical count or the Critical Customer Risk percentage criteria, the advisories are filtered to display only the critical customer risk Proactive Bug Notifications for the device.

From the Proactive Bug Notifications page, select an entry to view the **Device PBN Quick View** pane.

For more information on the fields on the Proactive Bug Notifications page and Device Quick View pane, see [Table 16 on page 61](#).

- **GENERAL**

Displays general information about the device.

For field descriptions, see [Table 17 on page 62](#).

- **CONTRACTS**

Displays information on the support contract for the device.

For field descriptions, see [Table 18 on page 63](#).

- **HARDWARE END OF LIFE DATES**

Displays information on the hardware End of Life (EOL) and End of Support (EOS) milestones for the device.

For field descriptions, see [Table 19 on page 63](#).

- **SOFTWARE END OF LIFE DATES**

Displays information on the software installed on the device.



NOTE: Software EOL information is available only for cloud-connected devices.

For field descriptions, see [Table 20 on page 63](#).

- Access support insights reports; No Link Title.

Field Description for Installed Base Page

Table 14: Fields on the Installed Base Tab


Field	Description
Model	The device model. For Unsupported devices, a red flag icon is displayed next to the device model. On hover, a pop-up appears indicating that the device is not recognized as being covered by an active support contract and may not be eligible for JTAC support.  If you find any data discrepancy, contact Customer Care for resolution. Once resolved, the data will be updated during the next data refresh cycle.

Table 14: Fields on the Installed Base Tab *(Continued)*

Field	Description
Serial Number	Unique ID mapped to the device.
Status	Service level to which the device is onboarded. Hover over the checkmark icon to view the states. - Blue checkmark—Indicates that the device is onboarded to the Assured service level. This service level includes access to the JSI experience. For more information, see "Integrate Your Juniper Support Resources to Your Organization" on page 24. - Grey checkmark—Indicates that the device is onboarded to the JSI service level only. NOTE: The field appears blank for devices not onboarded to the organization.
Device Host Name	Hostname of the device. NOTE: This field appears blank for devices not onboarded to the organization.
Action	Provides access to the device utilities.
End of Life	Hardware End of Life date for the device.
End of Sale	Hardware End of Sale date for the device.
End of Support	Hardware End of Support date for the device.
TSB/PSN	Reference number for a Technical Support Bulletin (TSB) or Product Support Notification (PSN) that communicates end of life notifications for the device type.
Software Release	Junos OS software version installed on the device.
Suggested Release	Recommended Junos OS software version for the device.
FRS	Date on which the Junos OS software version was first released.
EOE	Date on which the Junos OS software version installed on the device reaches end of engineering-level support.

Table 14: Fields on the Installed Base Tab *(Continued)*

Field	Description
EOS	Date on which the Junos OS software version installed on the device reaches end of support.

Table 15: Fields on the Security Incident Response Team Tickets Page of the *Device* Details Page

Field	Description
JSA ID	Unique value that identifies the security advisory on Juniper Networks Support Portal.
Title	Synopsis of the security advisory.
Severity	Severity rating of the security advisory. The values are: • Critical • High • Medium • Low
Affected Models	Device models affected by the security advisory.
OS Versions Affected	Junos or Junos Evo versions affected by the security advisory.
Release Date	Date on which the security advisory was first published.
JSA Updated Date	Date on which the security advisory was last updated.
Problem	Description of the security advisory.
Solution	Solution for the security vulnerability described in the advisory.
Workaround	Detailed explanation on how to temporarily resolve the problem.
Affected Series	Identifies one or more product series affected by the security advisory.
Release Notes	Short description of the security advisory.

Table 15: Fields on the Security Incident Response Team Tickets Page of the *Device* Details Page (Continued)

Field	Description
CVSS Score	Common Vulnerability Scoring System (CVSS) severity assessment score of the advisory in the range of 0-10. This field is available on the SIRT Quick View Pane only.
View SIRT details	Link to the advisory in the Juniper Networks Support Portal. You can view this link in the SIRT Quick View Pane .

Table 16: Fields on the Proactive Bug Notifications Page of the *Device* Details Page

Field	Description
ID	Unique value that identifies the Problem Report.
Headline	Synopsis of the problem.
Customer Risk	Classification of the potential impact to the customer if the bug was encountered in the network. The values include: • Critical—Conditions that could severely affect service, capacity or traffic, billing, and maintenance capabilities. • Major—Conditions that could seriously affect system operation, maintenance, administration, and so on. • Minor—Conditions that would not significantly impair the functioning of the network or significantly affect services.
Bug Type	Indicates the phase or activity during which the problem was discovered. Example: Day-1.
Trigger	Describes the events that happened before or at the time the problem occurred, or the event that caused the problem.
Introduced In	Junos or Junos Evo release where the problem was first found and reported.
Fixed In	Junos or Junos Evo release in which the problem was resolved.

Table 16: Fields on the Proactive Bug Notifications Page of the *Device Details Page (Continued)*

Field	Description
Release Notes	Short description of the problem.
Restoration	Indicates how the service can be restored when the problem occurs. Values include: • Self-recovery—Service, traffic, or operation disruptions are automatically restored without any user intervention. • Not-possible—It is not possible to restore the service or traffic. • Manual—User intervention is required to restore the service, traffic, or operation disruption.
Restoration Steps	Steps to restore the service when the problem occurs.
Workaround	Detailed explanation of how to temporarily resolve the problem until a permanent resolution is available.
Workaround Provided	Indicates whether a workaround for the problem is provided or not. Values include: • Yes—Workaround is available and is described in the Workaround field. • Not-possible—There are no workarounds to the problem.
Product Family	Identifies one or more products affected by the problem.

Table 17: Fields on the General Tile of the *Device Details Page*

Field	Description
Product SKU	Stock Keeping Unit (SKU) number assigned to the device.
Installed Address	Address of the site where the device is installed.
Software Version	Junos OS version installed on the device. NOTE: The field appears blank for non-cloud connected devices.

Table 17: Fields on the General Tile of the *Device Details Page (Continued)*

Field	Description
Recommended Release	Recommended Junos OS software version for the device.
Last Updated	Date on which the recommended Junos OS software was last updated.
Recommended Link	Link to a Juniper Support Portal Knowledge Base article with a list of recommended Junos OS versions for each Juniper platform.

Table 18: Fields on the Contract Tile of the *Device Details Page*

Field	Description
Service SKU	SKU assigned to the device's service contract.
Contract Type	Type of active support coverage provided for the device. Example: Maintenance.
Start Date	Date on which the service contract starts for the device.
End Date	Date on which the service contract ends for the device.
Reseller	Name of the reseller through which your company acquired the device.

Table 19: Fields on the Hardware End of Life Dates Tile of the *Device Details Page*

Field	Description
End of Life	Date on which the device reaches end of life.
End of Support	Date on which the device reaches end of support.

Table 20: Fields on Software End of Life Dates Tile of the *Device Details Page*

Field	Description
End of Support	Date on which the Junos OS software version installed on the device reaches end of support.

Table 20: Fields on Software End of Life Dates Tile of the *Device Details Page (Continued)*

Field	Description
End of Engineering	Date on which the Junos OS software version installed on the device reaches end of engineering-level support.
First Release Shipping	Date on which the Junos OS software version was first released.
Software EOL Link	Link to the Junos OS Dates & Milestones page in the Juniper support website. This page contains dates of important milestones for all Junos OS versions.

DC Edges and DC Switches

SUMMARY

This topic describes the DC Edges and DC Switches pages on the main menu.

IN THIS SECTION

- [Tasks You Can Perform | 65](#)
- [Field Description for DC Edges Page | 65](#)
- [Field Description for DC Switches Page | 66](#)

You can access information about the data center devices and the Apstra Edge devices in an organization on the Data Center Edges and Switches pages from the main menu on the left.

The **DC Edges** page displays the Apstra Edge devices and blueprints registered in the organization.

The **DC Switches** page displays detailed information about the devices in the data center managed by Apstra Data Center Director.



NOTE:

- You can filter the displayed devices on both pages using the **Org** drop-down—by individual sites, site groups, or all sites.

- If you are a superuser or a network admin with access to all sites, you can perform additional tasks on the DC Edges page.

Tasks You Can Perform

You can perform the following tasks on the **DC Edges** page:

- Adopt Apstra Edge devices—You can register an Apstra Edge Device by clicking **Adopt Apstra Edge** and complete the registration process.

See [Adopt Juniper Apstra Edge](#).

- View Apstra Edge devices registered in Juniper Data Center Assurance.
- Run Pre-flight Checks—You can run pre-flight checks to verify the configurations of the Edge device and the blueprints.

See [Pre-Flight Checks](#).

- Search by device name and view information about specific Apstra Edge devices.

You can perform the following tasks on the **DC Switches** page:

- View information about the devices in the data center managed by Apstra Data Center Director—You can view information about the devices managed by all the Apstra instances that are connected to Juniper Data Center Assurance.
- View information about specific devices by adding custom filters based on device name, vendor, device model, and so on.

Field Description for DC Edges Page

[Table 21 on page 65](#) lists the fields on the Data Centers tab.

Table 21: Fields on the DC Edges Page

Field	Description
Name	Name of the Apstra data center edge device.

Table 21: Fields on the DC Edges Page *(Continued)*

Field	Description
Registration Status	Indicates the registration status of Apstra data center edge device. The edge device installed in a data center managed by Apstra Data Center Director must be registered in Juniper Data Center Assurance application to receive data center event information from Apstra.
Management URL	URL of the Apstra Data Center Director instance.
Cloud Connectivity	Connectivity status of Juniper Data Center Assurance application with Juniper Apstra data center edge device. Juniper Data Center Assurance application can receive event information from Apstra Data Center Director only if it is connected with the edge device.
Apstra Connectivity	Connectivity status of Juniper Apstra Edge with Apstra Data Center Director.
Flow Server Status	Indicates whether the Apstra Flow Server is configured. Click Adopt Apstra Edge and enable Flow Server Details toggle button to configure Apstra Flow Server .

Field Description for DC Switches Page

Table 22 on page 66 lists the fields on the Switches tab.

Table 22: Fields on the DC Switches Page

Field	Description
Status	Status of the device managed by Apstra Data Center Director.
Name	Name of the device.
Vendor	Manufacturer of the device.

Table 22: Fields on the DC Switches Page *(Continued)*

Field	Description
OS Version	Version of operating system installed on the device.
MAC Address	MAC address assigned to the device.
MGMT IP	Management IP address assigned to the device.
Model	Device model; for example QFX5120-48Y.
Site	Site to which the device is assigned.
Serial Number	Serial number of the device.

RELATED DOCUMENTATION

[About the Inventory Page | 51](#)

[Adopt Juniper Apstra Edge | 124](#)

Validate Your Juniper Account Overview

User's identity is validated to prevent unauthorized access to customer-centric device information. Once validated, the user is granted access to the customer-centric information which the user is entitled to access.

To prevent unauthorized access, device information is classified into two categories, device-centric and customer-centric information. All users who have access to the organization can view device-centric details for a device. However, the user identity is validated before providing access to customer-centric details. This ensures that only authorized users have access to customer-centric information for a device.

[Table 23 on page 68](#) lists the device-centric and customer-centric details for a device.

Table 23: Device-centric and Customer-centric Details for a Device

	Installed Base Tab	Overview Tab	PBN Tab	SIRT Tab
Device-centric Information	- Assured - Product SKU - Serial Number - Model - HW EoL Date - HW EoS Date - Contract Type - Warranty Type	- Product SKU - Recommended Release - Last Updated - Recommended Link - Contract Type - End of Life - End of Support - End of Support - End of Engineering - First Release Shipping - Software EOL Link	- JSA ID - Title - Severity - CVSS Score - Affected Models - OS Versions Affected - Release Date - JSA Updated Date - Problem - Solution - Workaround - Affected Series - Release Notes	- ID - Headline - Customer Risk - Bug Type - Trigger - Introduced in - Fixed in - Release Notes - Restoration - Restoration Steps - Workaround - Workaround Provided - Product Family

Table 23: Device-centric and Customer-centric Details for a Device *(Continued)*

	Installed Base Tab	Overview Tab	PBN Tab	SIRT Tab
Customer-centric Information	• Installed Address • Contract ID • Contract SKU • Contract Start Date • Contract End • Ship Date • Reseller • Distributor • Warranty Start • Warranty End	• Installed Address • Contract SKU • Start Date • End Date • Reseller	-	-

For more information about validating your Juniper account, see ["Validate Your Juniper Account"](#) on [page 69](#).

Validate Your Juniper Account

SUMMARY

This topic describes the steps to validate your Juniper registered user account.

IN THIS SECTION

- [Validate Your Juniper Account | 71](#)
- [Validate Your Juniper Account as an Unregistered User | 72](#)
- [Users with Validated Juniper Account | 73](#)

Juniper Data Center Assurance validates users before providing access to customer-centric information about devices. The user credentials are validated through the Juniper Support Portal to prevent unauthorized access to device information. Once validated, the user is granted access to customer-centric information that the user is entitled to access. For more information on customer-centric information, see ["Validate Your Juniper Account Overview" on page 67](#).

Before you begin,

Usage Considerations

- If you are the user who successfully executed the ["Account Integration workflow" on page 24](#), your identity is automatically validated, and you can access the customer-centric JSI data artifacts associated with the linked account. If not, you must validate your identity to gain access to the customer-centric JSI data artifacts
- Validation is user-specific, each user must validate their identity individually to access the customer-centric data.
- Once validated, validation persists across user sessions, eliminating the need to re-validate your identity each time you log in.
- You can link multiple accounts to your organization. This allows you to get unified visibility into JSI artifacts associated with all of your linked accounts from the Installed Base.

When you link multiple accounts, the Installed Base shows customer-centric information only for the most recently linked account. The Installed Base hides customer-centric information from previously linked accounts. To access customer-centric information from an earlier linked account, you must re-validate that account.

For example, if you link Account 1 first and Account 2 later, Installed Base displays customer-centric information of devices associated with Account 2. Devices from Account 1 appears blurred. To access customer-centric information of devices associated with Account 1, you must re-validate it.

Prerequisites

You must fulfill the following conditions to gain access to the customer-centric information:

- You must have a registered Juniper account. Use this account to validate your identity.

You can register your account on the [Juniper Support Portal](#).

- You must have one of the following roles in JSP: Service, Service Direct, JSI Admin, JSI Standard, JSI Super Admin.

If you have any other role, your account validation will be unsuccessful.

- Your user account must have the required permission to access the information.

To have valid access privileges, your user account must be mapped to the Juniper account ID linked to the organization.

Validate Your Juniper Account

Customer-centric device information is masked to prevent unauthorized access. You must successfully validate your identity to access the masked customer-centric information of devices associated with the account, such as installed address, contract information, warranty information, and so. Once validated, validation persists across user sessions, eliminating the need for re-validation each time you log in.



NOTE: Validation is user-specific, each user must validate their identity individually to access the customer-centric data.

To validate your account:

1. Navigate to **Organization > Inventory > Installed Base** or **Organization > Support Insights > Assets & Contracts**.

On these pages, you can view device-centric information for all Juniper devices, but the customer-centric information for the devices appears blurred.

A banner appears prompting you to validate your user account. This is to confirm that you are authorized to view the customer-centric device information.

2. Do one of the following:

- Click **Validate Email ID** on the banner.
- Hover over any of the blurred fields and click **Validate**.

The Validate Juniper Account pop-up appears.

3. Enter your Juniper registered account credentials (email address and password) that you registered through [Juniper Support Portal](#).
4. Click **Validate**.

On successful validation, the page is refreshed and displays the device-centric information and the customer-centric information that you are entitled to access.



NOTE: Your access to customer-centric information depends on the access privileges granted to your account. For more information, see "[Users with Validated Juniper Account](#)" on page 73.

If the credentials you entered are recognized as invalid, you are prompted to re-enter valid credentials.

For further assistance, contact your service manager or customer support.

Validate Your Juniper Account as an Unregistered User

You must have a Juniper registered user account to access customer-centric device information.



NOTE: Validation is user-specific, each user must validate their identity individually to access the customer-centric data.

To validate your account as an unregistered user:

1. Navigate to **Organization > Inventory > Installed Base** or **Organization > Support Insights > Assets & Contracts**.

On these pages, you can view device-centric information for all the devices, but the customer-centric information appears blurred.

A banner appears, prompting you to validate your user account. This is to confirm that you are authorized to view the blurred customer-centric device information.

2. Do one of the following:


- Click **Validate Email ID** on the banner.
- Hover over any of the blurred fields and click **Validate**.

The Validate Juniper Account pop-up appears.

3. Enter your Juniper account credentials (email address and password).
4. Click **Validate**.

If the credentials you entered are recognized as invalid, you are prompted to re-enter valid credentials. You may also encounter this issue if you do not have a Juniper registered account. For further assistance, contact your service manager or customer support.

Figure 3: Invalid Credentials



Click **Close** to return to the page.

5. If you do not have a user account, click **Register** to register your account with Juniper.

You are redirected to the user registration page (<https://userregistration.juniper.net/>).

6. Register your account.

After successful registration, log in to the organization, navigate to the Installed Base tab or Assets & Contracts tab and repeat steps 2 through 4.

On successful validation, the page is refreshed and displays the device-centric information and the customer-centric information that you are entitled to access.



NOTE: Your access to sensitive customer-centric information depends on the access privileges granted to your account. For more information, see "[Users with Validated Juniper Account](#)" on page 73.

Users with Validated Juniper Account

This section describes the various levels of access controls you might encounter after you successfully validate your account.

All users who are invited to the organization can view device-centric information for all devices without any validation.

Users who successfully validate their credentials can view all, some, or none of the customer-centric information depending on whether they are authorized to access the same accounts to which the customer-centric information belongs.



NOTE: For assistance with access privileges, contact your service manager or Juniper customer support.

Table 24 on page 74 lists different user scenarios and the device-centric and customer-centric information accessible in each scenario.

Table 24: User Scenarios

User Scenario	Device-centric Data	Customer-centric Data
Validated User With All Matching Accounts	View all device-centric information	View customer-centric information of all Juniper devices linked to the organization. NOTE: When you link multiple accounts, the Installed Base shows customer-centric information only for the most recently validated account. The Installed Base hides customer-centric information from previously linked accounts. To view customer-centric information from an earlier linked account, you must re-validate that account.

Table 24: User Scenarios (Continued)

User Scenario	Device-centric Data	Customer-centric Data
Validated User With a Few Matching Accounts	View all device-centric information	View customer-centric information of only the Juniper devices you are entitled to access. NOTE: When you link multiple accounts, the Installed Base shows customer-centric information only for the most recently validated account. The Installed Base hides customer-centric information from previously linked accounts. To view customer-centric information from an earlier linked account, you must re-validate that account.
Validated User With No Matching Accounts	View all device-centric information	No access to customer-centric information of Juniper devices linked to the organization

Configure Ports

IN THIS SECTION

- [Configure Firewall to Access AWS S3 Buckets | 78](#)

DC Assurance requires specific outbound and inbound network connectivity from your Apstra VM and Flow server to cloud services and internal systems. Your network and security teams must configure firewall rules to permit traffic on these ports and to the required endpoints.

[Table 25 on page 76](#) lists the ports and endpoints that should be enabled on the firewall.

Table 25: Ports Configuration

Direction	Service Type	Portal	Port Number	Description
Outbound	DC Assurance Admin Portal	dc.ai.juniper.net	TCP 443	Opens the DC Assurance GUI.
Outbound	Apstra Edge to Juniper Cloud Connectivity	ep-term.ai.juniper.net	TCP 443	Establishes a connection from the Apstra Edge to a cloud termination service for secure communication with cloud infrastructure.
Outbound	Apstra Edge to Amazon S3 Buckets	- usw2-ai-prod-dc-payloads.s3.us-west-2.amazonaws.com - usw2-ai-prod-apstra-edge-repo.s3.us-west-2.amazonaws.com	TCP 443	Apstra Edge interacts with Amazon S3 buckets to: - download edge binaries during remote upgrade from JCloud. - upload telemetry data. See "Configure Firewall to Access AWS S3 Buckets" on page 78 below for more information.
Outbound	Apstra Edge to Apstra Controller	<i>IP address of the Apstra instance</i>	TCP 443	Performs configurations and retrieves data such as blueprints and anomalies from the Apstra controller.

Table 25: Ports Configuration *(Continued)*

Direction	Service Type	Portal	Port Number	Description
Inbound	Apstra Anomalies Stream Receivers to Apstra Edge	localhost	TCP 9595	Receives pushed anomalies data from Apstra stream receivers.
Inbound	Apstra Metrics Stream Receivers to Apstra Edge	localhost	TCP 9797	Receives pushed metrics data from Apstra stream receivers.
Inbound	Local Admin Server to Apstra Edge	- localhost 0.0.0.0	TCP 8081	Used for internal health checks, metrics, and REST API access.
Outbound	Apstra Edge to Flow Server	FlowServer OpenSearch endpoint	TCP 9200	Queries OpenSearch for flow and analytics data.
Outbound	Apstra Edge to Flow Server	Metrics endpoint	TCP 8080	Retrieves metrics exposed by the Controller.
Outbound	Apstra Edge to vCenter Server API	vCenter Server API	TCP 443	vSphere API (SOAP/SDK) is used by VMware SDK clients to authenticate, and query inventory vCenter.
Outbound	Apstra Edge to Slurm	<i>IP address of the Slurm REST API Server instance</i>	TCP 6820	Queries the Slurm REST API server periodically to fetch workload/job information.

Configure Firewall to Access AWS S3 Buckets

Configure FQDN-Based Firewall Rules (Recommended)

Follow this approach if your firewall supports FQDN / DNS-based rules.

AWS S3 uses dynamic IP addresses, and the FQDN rules automatically track IP address changes. This allows you to restrict access to specific buckets only and does not require ongoing IP maintenance.

What to allow—Outbound HTTPS (TCP 443) from the Edge Service to the specific S3 bucket hostnames in the following format: *bucket-name.s3.region.amazonaws.com*

Configure IP-Based Firewall Rules

If your firewall does not support FQDN-based rules, you may allow access using AWS-published IP ranges.

AWS S3 IP ranges are not static and may change at any time. Therefore, you must keep your firewall rules up to date to avoid service disruption. The Edge Service must also be able to resolve AWS endpoints via DNS. You must ensure that outbound access to your DNS resolver(s) is permitted.

What to allow—Outbound HTTPS (TCP 443) from the Edge Service to the destination IP ranges corresponding to the S3 AWS service and the AWS region(s) used by the service.

For the complete list of AWS regions and IP address ranges, see <https://ip-ranges.amazonaws.com/ip-ranges.json>.



NOTE: To refine your search for the relevant region and IP ranges from the complete list, filter for entries with:

- "service": "S3"
- "region": "*applicable-region*"

6

CHAPTER

Contracts Management

IN THIS CHAPTER

- [About the Contracts Page | 80](#)
-

About the Contracts Page

IN THIS SECTION

- [Stats Panel | 80](#)
- [Contracts Table | 83](#)
- [Tasks You Can Perform | 86](#)

The Contracts dashboard (**Organization > Admin > Contracts**) displays the support contract details for all Juniper devices linked to your organization, whether onboarded directly or through the account integration workflow.

Use this dashboard to view support details, monitor support coverage, identify unsupported devices, and prioritize contract renewals.



NOTE:

- You can view contract information only for your onboarded devices.

You can onboard devices directly or through the ["account integration workflow" on page 88](#).

- If your account is not validated, all customer-centric information in the dashboard appears blurred. Validate your account to view the information, see ["Validate Your Juniper Account" on page 86](#).


The Contracts dashboard is organized into two main sections: Stats Panel and Contracts Table.



NOTE: The dashboard is refreshed every 24 hours.

Stats Panel

The stats panel provides a high-level visual summary of contract coverage across your organization. Each category also acts as an interactive filter for the Contracts table.



- **Device Model**—Displays the number of devices in the organization per model.

Use this category to identify distribution of device models in your organization.

Click a model to filter the table to display only those devices.

- **Contract Status**—Displays the percentage of devices in each support contract status category.

Table 26: Contract Status Category

Status	Description
Declined	Support contract is not purchased for the device.
EoS	Support contract for the device has reached End of Support (EoS).
Service Available	Support contract is available on purchase for the device.
Active	Support contract for the device is active.

Click a category to filter the table based on the selected contract status.

- **Contract Type**—Displays the percentage of devices associated with each support service type available to the customer.

Table 27: Contract Type

Field	Description
SVC	Support service provided directly by Juniper.
PAR	Support service provided through an authorized Juniper partner.
Maintenance	Active support service present for the device.

Click a contract type to filter the table accordingly.

- **Contract State**—Displays the percentage of devices that are currently Supported or Unsupported.
 - **Supported**—Devices covered by an active support contract in the Juniper support database. These devices are eligible for JTAC support.
 - **Unsupported**—Devices not covered by an active support contract in the Juniper support database. JTAC support may not be available for these devices until the contract is purchased or renewed.

If you notice any data discrepancy, contact Customer Care for resolution. Once resolved, the data will be updated during the next data refresh cycle.

Click a contract state to filter the table by that state.



NOTE:

- By default, the Supported filter is selected. You can click it again to clear the filter.
- Unsupported devices can still be onboarded and managed within your organization.

- **Contract Expiration**—Displays the percentage of devices grouped by contract expiration timelines.

Table 28: Expiration Timelines

Time Range	Description
Expired	Indicates the percentage of devices with an expired contract.
0 - 3 months	Indicates the percentage of devices whose contracts are expiring within the next 3 months.
0 - 12 months	Indicates the percentage of devices whose contracts are expiring within the next 12 months.
> 12 months	Indicates the percentage of contracts whose expiry is after 12 months.

Click a category to filter the table based on contract expiration timeline.

Contracts Table

The Contracts table provides detailed contract insights for the devices, allowing you to perform deeper analysis and troubleshooting.

You can apply filters from the Stats Panel or perform keyword-based searches to refine the table records.



TIP: Use the Stats Panel to quickly identify areas of concern, then explore the Contracts Table for further investigation.

By default, the Contracts dashboard applies the **Supported** filter, so only devices with active support contracts appear when the page is opened.

To view all devices regardless of their contract status, clear the **Supported** filter in the Stats Panel.

For field description, see [Table 29 on page 84](#).



NOTE: If your account is not validated, all customer-centric information in the table appears blurred. Validate your account to view the information, see "[Validate Your Juniper Account](#)" on page 86.

Table 29: Field Description


Field	Description
Model	The device model. For Unsupported devices, a red flag icon is displayed next to the device model. On hover, a popover appears indicating that the device is not recognized as being covered by an active support contract and may not be eligible for JTAC support.  Please contact Customer Care at 888-314-5822 as your device does not appear to be covered by an Active Support contract which may lead to a denial of service should you have a need to contact J-Tac If you notice any data discrepancy, contact Customer Care for resolution. Once resolved, the data will be updated during the next data refresh cycle.
Serial Number	Unique ID mapped to the device.
Status	Service level to which the device is onboarded. - Blue checkmark—Indicates that the device is onboarded to the Assured service level. This service level includes access to the JSI experience. - Grey checkmark—Indicates that the device is onboarded to the JSI service level only. NOTE: This field appears blank for devices not onboarded to the organization.
Device Host Name	Hostname of the device. NOTE: This field appears blank for devices not onboarded to the organization.
Installed Address	Location where the device is installed.

Table 29: Field Description (*Continued*)

Field	Description
Service SKU	Stock Keeping Unit (SKU) number assigned to the device's support contract. Types of SKUs: • SVC-xxx—SKU for the support service provided directly by Juniper to the customer. • PAR-xxx—SKU for the support service provided by an authorized Juniper partner to the customer.
Contract ID	Support contract number assigned to the device.
Support Contracts Status	Status of the support contract for the device. • Declined—Support contract is not purchased for the device. • EoS—Support contract for the device has reached EOS. • Service Available—Support contract is available on purchase for the device. • Active—The device has an active support contract.
Contract Type	Type of support service provided to the customer for the device. • SVC—Support service provided directly by Juniper. • PAR—Support service provided through an authorized Juniper partner. • Maintenance—Active support service present for the device.
Contract Start Date	Support contract start date for the device.
Contract End Date	Support contract end date for the device.
Warranty Type	Warranty type associated with the device. Example: Standard Hardware Warranty.
Warranty Start Date	Start date of warranty for the device.
Warranty End Date	End date of warranty for the device.

Table 29: Field Description (*Continued*)

Field	Description
Reseller	Juniper authorized reseller of the device.
Distributor	Juniper authorized distributor of the device.

Tasks You Can Perform

You can perform the following tasks on the Contracts dashboard:

- **Validate your account**

To view customer-centric information for the device, you must validate your account. Customer-centric information on the page appears blurred if your account is not validated.

To validate your identity and access the customer-centric information, do one of the following:

1. Do one of the following

- Click **Validate Email ID** on the banner.
- Hover over any of the blurred fields and click **Validate**.

2. Enter your Juniper support credentials (created through the [Juniper Support Portal](#)) and click **Validate**.

Once validated, you can view the customer-centric information you are authorized to access.

If the system finds the credentials you entered to be invalid, you are prompted to re-enter valid credentials.

You may also encounter this issue if you do not have a Juniper registered account. You can register your account with Juniper through the [Juniper Support Portal](#). For further assistance, contact your Service Manager or Juniper customer care.

- **Find a device using filters**

The Contracts dashboard offers multiple ways to narrow down your search.

- **Use Stats Panel filters**—The Stats Panel supports interactive filters.

Click a category to show only matching devices in the table. To clear a filter, click the button again.

For example, Click **MX480** to view contract insights for only MX480 devices in the table.

- **Apply batch filters**—Combine multiple filters to narrow down your search. Batch filters let you apply more than one criterion at the same time.

For example, select SRX4200 from Device Model and Unsupported from Contract State to show only unsupported SRX4200 devices in the table.



NOTE: Active filters appear grayed out on the Stats Panel. To clear a filter, click the selected item again.

- **Enter keywords in the Filter box**—Start typing in the **Filter** box.

For example, start typing the model name. As you type each character, matching entries appear in the drop-down list. Click the device that you want to show in the table. To clear the filter, click **Clear All**.

- **Modify table settings**

Click the column settings icon on the top-right corner to show, hide, and reorder the columns so that the table shows exactly the information that you want to see.



In the Table Settings window, you can:

- Select or clear column checkboxes
 - Drag columns to reorder them
 - Use the search box to quickly locate a column
 - Use **Select All** or **Deselect All** as needed
 - Click **Close** to apply changes
- **Export contract details in CSV format**

To export details of all devices, click downloads icon at the top-right corner. The details are exported to a CSV file that you can download to your local system.



NOTE: Any non-English characters might appear as special characters when you open the file. To prevent this issue, follow these steps:

1. Open a new Excel file and then select **File > Import > CSV File > Import**.
2. Select the file to be opened and then click **Get Data**.
3. In the Text Import Wizard, select **Unicode (UTF-8)** as the **File Origin**.
4. Click **Finish**.

- **Integrate Your Juniper Support Resources with Your Organization**

To integrate devices associated with your Juniper-registered account with your organization, you must link your Juniper-registered credentials (created through the [Juniper Support Portal](#)) to your organization.

After the account is successfully linked to the organization, you can access support insights for the devices from the Installed Base. This includes all primary and secondary assets associated with your registered email address that is maintained in Juniper's support database and covered by an active Juniper support contract

You can link multiple accounts with the organization.

To integrate your Juniper-registered account with the organization:

1. Click **Add Juniper Account**.

The Add Juniper Account pop-up appears.

2. Enter your Juniper support credentials (created through the [Juniper Support Portal](#)).
3. Click **Add**.

On successful integration, the primary account name associated with your credentials is listed in the Linked Accounts section. The Linked Account section also lists other primary accounts that are currently linked to the organization.

Alternatively, you can link your account to the organization from the Settings page (**Organization > Settings > Juniper Account Integration**). If you no longer wish to view data insights for devices associated with an account in your organization, you can unlink the account from this page. This removes the devices associated with the account from the organization. For more information, see ["Integrate Your Juniper Support Resources to Your Organization" on page 24](#).

- **Access the Installed Base**

Click **Installed Base Inventory** to navigate to the Installed Base page. From this page, you can onboard devices to cloud and view device-specific support insights, including hardware and software milestones, security vulnerabilities, and bug notifications.

7

CHAPTER

Advisories

IN THIS CHAPTER

- [About the Advisories Page | 90](#)
-

About the Advisories Page

IN THIS SECTION

- [SIRTS Dashboard | 90](#)
- [PBN Dashboard | 93](#)
- [Tasks You Can Perform | 96](#)

The Advisories (**Organization > Advisories**) page provides a centralized view of security vulnerabilities and known software issues affecting the devices in your organization. It helps you monitor risks, assess their impact, and take appropriate remedial actions.

The page contains two tabs: "[SIRT](#)" on page 90 (Security Incident Response Team) and "[PBN](#)" on page 93 (Proactive Bug Notifications).

Use the interactive dashboards on these tabs to analyze affected devices, prioritize remediation effort based on severity, and export insights for offline use or further customization. By using the Advisories page, you can stay informed about the vulnerabilities impacting your onboarded Juniper Network devices.

SIRTS Dashboard

The Juniper Networks Security Incident Response Team (SIRT) issues advisories to inform you about security vulnerabilities and available fixes for Juniper Networks products and services.

The **SIRT** dashboard provides a summary report on the security vulnerabilities advisories published by the SIRT team for the Juniper Networks devices onboarded to the organization.

Use the dashboard to view the list to identify affected devices, assess the impact and severity of the vulnerabilities, view recommended actions to mitigate the risk, and so on.

To access the SIRT dashboard, navigate to **Organization > Advisories > SIRT**.



NOTE: The advisories are available only for cloud-connected devices linked to your organization.


The [SIRT insights panel on page 91](#) provides you an overview of the security vulnerabilities impacting your organization's onboarded devices.

This interactive panel helps you identify which devices are at risk, prioritize security patching, understand where the risk is concentrated, and take targeted action.

Use the different filters to pivot the report and get focused insights into the vulnerabilities impacting your devices.

You can filter data based on the following parameters:

Figure 4: SIRT Insights Panel



- **Product**— Displays the device models in your inventory that are affected by security vulnerabilities. Helps you identify which device models are the most affected by vulnerabilities.

The row displays the top three device models with the highest number of SIRTs. Remaining models are grouped under **Others**.

Entries are sorted in descending order based on the number of devices per model affected by security vulnerabilities. Each entry shows device count (in units and percentage of total) and number of associated SIRTs for each model.

Click a model to filter security vulnerabilities specific to that model.

- **SW Release**—Displays the Junos OS release version running on the affected devices. Helps you track which OS version is running on the affected devices.

The row displays the top OS releases with the highest number of SIRTs. Remaining releases are grouped under **Others**.

Entries are sorted in descending order based on the number of devices per release with associated security vulnerabilities. Each entry shows the device count (in units and percentage of total) and number of associated SIRTs.

Click any release to filter SIRTs advisories specific to that version of Junos OS.

- **SIRTs by Age**—Displays the distribution of SIRTs based on their publication age. This helps you assess how long vulnerabilities have existed and whether older advisories are still impacting your devices.

Select a time range to filter the advisories published within that period.

- **Exposure by Severity**—Displays the percentage of security vulnerabilities impacting the devices grouped by severity: Critical, High, Medium, and Low. The percentage indicates how much of your inventory is affected by each severity level.

This helps you filter vulnerabilities by severity, identify what percentage of your inventory is impacted by vulnerabilities, and focus on remediating the most critical vulnerabilities.

Click a severity level to filter advisories based on that level.

To review detailed SIRT records and identify affected devices and advisories, use the SIRT table. For field description, see [Table 30 on page 92](#).

Table 30: Fields on the SIRT Dashboard

Field	Description
ID	Unique value that identifies the security advisory on Juniper Networks' Support Portal.
Title	Synopsis of the security advisory.
Affected Devices	Devices impacted by the security advisory. Click the hyperlink to view details of the devices impacted by the security advisory. The details include device serial number, model, cloud-connected status, hostname, and the Junos OS version installed on the device. You can also export the details as a CSV file for offline access or further analysis.
Severity	Severity rating of the security advisory. The values are: • Critical • High • Medium • Low
CVSS Score	Common Vulnerability Scoring System (CVSS) is the severity assessment score of the advisory in the range of 0-10. A score of 0 indicates low severity and a score of 10 indicates critical severity.
Affected Models	Device models impacted by the security advisory.
OS Versions Affected	Junos or Junos Evo versions impacted by the security advisory.

Table 30: Fields on the SIRT Dashboard *(Continued)*

Field	Description
Published Date	Date on which the security advisory was first published.
Updated Date	Date on which the security advisory was last updated.
Problem	Description of the security advisory.
Solution	Solution for the security vulnerability described in the advisory.
Workaround	Detailed explanation on how to temporarily resolve the problem.
Release Notes	Additional details related to the the security advisory, such as links to relevant KB articles.

Click a SIRT advisory to view detailed information about the advisory in a side panel. From the panel, you can click **View SIRT Details** button to navigate to the security bulletin in the Juniper Support Portal.

PBN Dashboard

Juniper Networks notifies you about known software issues through Proactive Bug Notifications (PBNs). Each PBN includes information such as bug description, impact level, affected devices, and the software releases in which the issue is resolved.

The **PBN** dashboard displays a summary report on the known software issues that impact Juniper Networks devices onboarded to the organization.

Use the dashboard to view the list of devices impacted by the issue, severity rating assigned to the issue, its potential impact, possible fixes or workarounds to address the issue, and so on.

To access the PBN dashboard, navigate to **Organization > Advisories > PBN**.




NOTE: The report is available only for cloud-connected devices linked to your organization.

The [PBN insights panel on page 94](#) provides you an overview of the known issues impacting your organization's onboarded devices. This interactive panel helps you identify high-risk devices, understand the scope of impact, prioritize troubleshooting, and schedule maintenance and upgrades.

Use the different filters to pivot the report and perform targeted investigation into the issues impacting your devices.

You can filter data based on the following parameters:

Figure 5: PBN Insights Panel



- **Product**—Displays the device models in your inventory most affected by known issues. Helps you identify which device models are the most affected by known issues.

The row displays the top three device models with the highest number of known issues. Remaining models are grouped under **Others**.

Entries are sorted in descending order based on the number of affected devices per model. Each model shows model's device count (in units and percentage of total) and number of associated PBNs.

Click a model to filter known issues affecting that model.

- **SW Release**—Displays the Junos OS release versions running on the affected devices. Helps you track which OS version is running on the affected devices.

The row displays the top OS releases with the highest number of known issues. Remaining models are grouped under **Others**.

Entries are sorted in descending order based on the number of affected devices per release. Each version shows the device count (in units and percentage of total) and the number of associated PBNs.

Click a release to filter known issues by that specific Junos OS version.

- **Exposure by Customer Risk**—Lists the percentage of known issues grouped by severity: Critical, Major, and Minor.

This helps you sort known issues by severity, identify what percentage of your inventory is impacted by the issues, and plan upgrades.

Click a severity level to filter PBNs based on that rating.

To review detailed bug notifications and identify affected devices, issues, and remediation information, use the PBN table. For field description, see [Table 31 on page 95](#).

Table 31: Fields on the PBN Dashboard

Field	Description
ID	Unique value that identifies the Problem Report.
Title	Synopsis of the problem.
Affected Devices	Devices impacted by the problem. Click the hyperlink to view details of the devices impacted by the known issue. The details include device serial number, model, cloud-connected status, hostname, and the Junos OS version installed on the device. You can also export the details as a CSV file for offline access or further analysis.
Customer Risk	Classification of the potential impact to the customer if the bug is encountered in the network. The values include: • Critical—Conditions that could severely affect service, capacity or traffic, billing, and maintenance capabilities, and require immediate corrective action. • Major—Conditions that could seriously affect system operation, maintenance, administration, and require attention. • Minor—Conditions that would not significantly impair the functioning of the network or significantly affect services.
Bug Type	Indicates the phase or activity during which the problem was discovered. Example: Day-1, Regression.
Affected Model	Device models impacted by the problem.
OS Versions Affected	Junos or Junos Evo versions impacted by the problem.
Introduced In	Junos OS or Junos OS Evo release in which the problem was first found and reported.
Fixed In	Junos OS or Junos OS Evo release in which the problem was resolved.
Release Notes	Additional details on the problem.

Table 31: Fields on the PBN Dashboard (*Continued*)

Field	Description
Restoration	Indicates how the service can be restored when the problem occurs. Values include: • Self-recovery—Service, traffic, or operation disruptions are automatically restored without any user intervention. • Not-possible—It is not possible to restore the service or traffic. • Manual—User intervention is required to restore the service, traffic, or operation disruption.
Workaround	Detailed explanation of how to temporarily resolve the problem until a permanent resolution is available.
Workaround Provided	Indicates whether a workaround for the problem is provided or not. Values include: • Yes—Workaround is available and is described in the Workaround field. • Not-possible—There are no workarounds to the problem.

Click a PBN to view detailed information about the issue in a side panel.

Tasks You Can Perform

You can perform the following tasks from this page:

- Modify table settings—Click the column settings icon on the top-right corner to show, hide, and reorder the columns so that the table shows exactly the information that you want to see.



In the Table Settings window, you can:

- Select or clear column checkboxes
- Drag columns to reorder them
- Use the search box to quickly locate a column
- Use **Select All** or **Deselect All** as needed

- Click **Close** to apply changes
- Access the Installed Base inventory—Click **Installed Base Inventory** to navigate to the Installed Base page. From this page, you can onboard devices to cloud and view device-specific support insights, including security vulnerabilities and bug notifications.
- Export details of all the devices in CSV format—To export details of all devices, on the respective tab, click **Download CSV** at the top-right corner. The details are exported to a CSV file that you can download to your local system.



NOTE: If you open the downloaded CSV file with Microsoft Excel on a Mac computer, any non-English characters in the file might appear as special characters. To prevent this issue, follow the steps below:

1. Open a new Excel file and then select **File > Import > CSV File > Import**.
2. Select the file to be opened and then click **Get Data**.

The **Text Import Wizard** window appears.
3. Select **Unicode (UTF-8)** as **File Origin**.
4. Click **Finish**.

- Enter the keyword in the search box to filter data based on keyword.

8

CHAPTER

Support Insights

IN THIS CHAPTER

- [About the Support Insights Page | 99](#)
-

About the Support Insights Page

IN THIS SECTION

- [Alarms Dashboard | 99](#)
- [Tasks You Can Perform | 102](#)

The Support Insights page provides dashboards that deliver data-driven insights into the operational health of devices in your organization. These insights help you monitor issues, understand device status, and take preventive or corrective actions.

To access Support Insights, navigate to **Organization > Support Insights** or **Organization > Inventory > Installed Base > Support Insights**.

If you want to access device-specific insights, navigate to the Installed Base (**Organization > Inventory > Installed Base**) page.



NOTE:

- To view data in the support insights dashboard, ensure that at least one Juniper account is linked to the organization or there are onboarded devices in the organization. Otherwise, the dashboard will appear blank.
- Support insights for an onboarded device are available only in the organization to which it is onboarded.

The following sections describe the dashboards available on the Support Insights page.

Currently, the Support Insights page supports only the Alarms dashboard.



NOTE: The dashboards are refreshed every 24 hours.

Alarms Dashboard

Devices trigger alarms when an abnormal event prevents them from functioning optimally.

The Alarms dashboard provides detailed information on the alarms generated on the devices onboarded to the organization.


Use the dashboard to:

- Monitor alarms generated on onboarded devices
- Identify issues based on severity and device model
- Prioritize remediation actions to minimize downtime
- Export dashboard data as CSV for offline analysis

To access the Alarms dashboard, navigate to **Organization > Support Insights > Device Health > Alarms**.

 **NOTE:** Alarms generated by onboarded devices appear in the Alarms dashboard after 24 hours.

You can filter data based on the following parameters:



- The total number of devices grouped by device model.

For example, MX240-AC and MX240-DC is grouped under MX240.

Click a device model to filter data for that device model.

- You can filter data based on alarm severity.

The severity levels are:

- Major—Indicates a critical condition on the device and requires immediate action to prevent further service disruption.
- Minor—Indicates a non-critical condition on the device but requires monitoring or maintenance to prevent potential service disruption.

Click a severity level to filter alarms based on the severity.

Use the Alarm table for detailed report on the Alarms. For field descriptions, see [Table 32 on page 101](#).

Table 32: Fields on the Alarms Table

Field	Description
Model	Device model.
Serial Number	Unique ID mapped to the device.
Assured	Service level to which the device is onboarded. Hover over the checkmark icon to view the states. - Blue checkmark—Indicates that the device is onboarded to the Assured service level. This service level includes access to the JSI experience. - Grey checkmark—Indicates that the device is onboarded to the JSI service level only. NOTE: The field appears blank for devices not onboarded to the organization.
Device Host Name	Hostname of the device on which the alarm occurred.
Alarm Status	Status of the issue that raised the alarm. - Active—Alarm is still active and requires attention. - Clear—Alarm is not active. It is either fixed or closed by the device.
Last Reported	Date when the alarm was last raised.
Alarm Time Local	The local time at which the alarm was recorded. The time is adjusted to the local time zone of the system or the user viewing the report.

Table 32: Fields on the Alarms Table *(Continued)*

Field	Description
Class	The severity level of the alarm. The options are: • Major—Indicates a critical condition on the device and requires immediate action to prevent further service disruption. • Minor—Indicates a non-critical condition on the device and requires monitoring or maintenance to prevent potential service disruption.
Description	A short description of the issue that triggered the alarm.

Tasks You Can Perform

You can perform the following tasks from this page:

- Modify table settings—Click the column settings icon on the top-right corner to show, hide, and reorder the columns so that the table shows exactly the information that you want to see.



In the Table Settings window, you can:

- Select or clear column checkboxes
- Drag columns to reorder them
- Use the search box to quickly locate a column
- Use **Select All** or **Deselect All** as needed
- Click **Close** to apply changes
- Access the Installed Base inventory—Click **Installed Base Inventory** to navigate to the Installed Base page. From this page, you can onboard devices to cloud and view device-specific support insights, including security vulnerabilities and bug notifications.
- Export details of all the devices in CSV format—To export details of all devices, on the respective tab, click **Download CSV** at the top-right corner. The details are exported to a CSV file that you can download to your local system.



NOTE: If you open the downloaded CSV file with Microsoft Excel on a Mac computer, any non-English characters in the file might appear as special characters. To prevent this issue, follow the steps below:

1. Open a new Excel file and then select **File > Import > CSV File > Import**.
2. Select the file to be opened and then click **Get Data**.

The **Text Import Wizard** window appears.

3. Select **Unicode (UTF-8)** as **File Origin**.
4. Click **Finish**.

- Enter the keyword in the search box to filter data based on keyword.

9

CHAPTER

Audit Logs

IN THIS CHAPTER

- [Audit Logs Overview | 105](#)
 - [About the Audit Logs Page | 105](#)
-

Audit Logs Overview

An audit log is a record of activities initiated by users. You can view a record of user-initiated activities such as accessing, creating, updating, or deleting any resource or component. Audit logs are useful in tracking and maintaining a history of these activities.



NOTE: Audit logging does not track device-initiated activities. Audit logs are cleared every 30 days.

Superusers and network administrators can view and filter audit logs to determine which users performed which actions at what time.

For example, a super user or network administrator can use audit logs to see who:

- added user accounts on a specific date.
- accessed the organization and at what time.
- added or deleted a site.

RELATED DOCUMENTATION

[About the Audit Logs Page | 105](#)

About the Audit Logs Page

IN THIS SECTION

- [Tasks You Can Perform | 106](#)
- [Field Descriptions | 106](#)

To access this page, select **Organization > Audit Logs**. Superusers and network administrators can view and filter audit logs for the organization. The Audit Logs page refreshes automatically and displays the latest logs.

Tasks You Can Perform

- View details of an audit log—Select an audit log to view basic information about the log, such as timestamp, the name of the user who initiated the task, the log message, and site details.

The View Details field appears for logs that are generated based on modifications made by the user to existing settings. Click **View Details** to view details about the update.

- Filter the data displayed in the table—You can filter the audit logs based on time interval, administrators, sites, and log messages.
 - To filter logs based on time interval you select, click *Time Interval* (by default it is set to Today). You can choose Last 60 Minutes, Last 24 Hours, Last 7 Days, Custom Date (enter a custom date) Today, Yesterday, This Week, or Custom (enter a custom time range).
 - To filter logs based on administrators, click **All Admins** (by default it is set to All Admins) or on the downward arrow, from the drop-down list, select the check boxes next to the administrators' names to filter logs by their activity.
 - To filter logs based on sites, click **All Sites** (by default it is set to All Sites) or on the downward arrow, from the drop-down list, select the check boxes next to the sites to display logs for the selected sites.
 - To filter logs based on messages, in the **Search by message** field, enter one or more keywords.

To remove the administrators and sites filter criteria, clear the selected check boxes in the drop-down list. To remove the messages filter criteria, click **Close** (x) in the check box.

Field Descriptions

Table 33 on page 106 describes the fields on the Audit Logs page.

Table 33: Fields on the Audit Logs Page

Field	Description
Timestamp	Date and time at which the audit log was recorded.
Admin Name	Name and e-mail address of the user who initiated the task.
Message	Description of the logged task.
Site	Name of the site in which the task was initiated.

Table 33: Fields on the Audit Logs Page *(Continued)*

Field	Description
View Details	Displays a clickable View Details link if any modifications were made by the user to existing settings, such as updating user roles, site settings, and so on.

RELATED DOCUMENTATION

| [Audit Logs Overview](#) | 105

10

CHAPTER

Licensing

IN THIS CHAPTER

- [Juniper Data Center Assurance Licensing and Subscriptions | 109](#)
-

Juniper Data Center Assurance Licensing and Subscriptions

IN THIS SECTION

- [Trial Period and Grace Period](#) | 110

You must have a valid subscription with the appropriate license tier to use features in Juniper Data Center Assurance. This topic describes the available license tiers and the features included in each tier.

Licenses for Juniper Data Center Assurance are derived from Apstra Data Center Director subscriptions. The license tier (Standard, Advanced, or Premium) assigned in Apstra determines the features available in Data Center Assurance.

Apstra licenses follow a subscription-based model, where each subscription defines:

- The license tier (Standard, Advanced, or Premium)
- The subscription duration (1, 3, 5, or 7 years)
- The number of licenses purchased

Feature access in Data Center Assurance is provided through entitlements, which are derived from your active subscriptions.

For more information about Apstra Data Center Director licenses, see [Software Licenses for Apstra Data Center Director](#).

To view your subscriptions and access associated entitlement (licensed features) in Juniper Data Center Assurance, you must link your Juniper Networks account in the **Juniper Account Integration** tile on the Organization Settings (**Organization** > **Settings**) page. For more information about linking your Juniper Networks account, see ["Integrate Your Juniper Support Resources to Your Organization" on page 24](#).

After linking your account, administrators can view subscriptions associated with their Apstra Data Center Director on the Subscriptions page (**Organization** > **Subscriptions**).

The Subscription page displays subscription type, entitlement quantity, subscription status, and validity period.

If you attempt to access a feature that is not included in your license tier, Juniper Data Center Assurance displays a warning message indicating that a valid subscription is required. To obtain access, contact your designated HPE Networking representative.

[Table 34 on page 110](#) lists license tiers and the features available in each.

Table 34: Licenses

License Tiers	Features
Standard	• Marvis AI Assistant • Dashboard • Alerts
Advanced	• Marvis AI Assistant • Dashboard • Alerts
Premium	• Marvis AI Assistant • App Aware • Impact Analysis • Dashboard • Alerts • Predictive Analytics • Network Sustainability Analytics

Trial Period and Grace Period

New users can access the premium features in Juniper Data Center Assurance for a 90-day trial period starting from the date of onboarding Apstra Edge.

After the trial period ends, a valid subscription is required to continue using premium features.

Existing licenses have a 30-day grace period from the date of expiry. After the grace period ends, your subscription expires and the associated feature entitlements are no longer available.

To purchase or renew subscriptions, contact your designated HPE Networking representative.

11

CHAPTER

Marvis AI Assistant for Data Center

IN THIS CHAPTER

- [Marvis AI Assistant for Data Center | 113](#)
 - [About the Marvis Page | 113](#)
 - [View Data Center Events in Marvis Actions | 114](#)
 - [Event Types Displayed in Marvis Actions | 115](#)
 - [Marvis AI Assistant Overview | 118](#)
 - [Find Information and Troubleshoot Juniper Products and Services Using Marvis AI Assistant | 119](#)
-

Marvis AI Assistant for Data Center

A key component of Juniper Data Center Assurance is the Marvis AI Assistant for Data Center, which is the AI-native virtual network assistant for data center operations. Marvis AI Assistant for Data Center combines the power of AI and the intent-driven approach of Apstra in building and operating data center.



Video: [Marvis AI Assistant for Data Center](#)

Marvis AI Assistant ingests the network event information received from the Apstra Edge component, analyzes them and provides network administrators actionable insights and recommendations to resolve an event. Marvis AI Assistant performs root cause analysis for the events received from the edge and displays them in Marvis Actions under different event categories. Administrators can quickly examine such events at the click of a button or, if more troubleshooting is required, they can access Apstra Data Center Director and perform the recommended steps to resolve the issue. All these can happen in real-time before network traffic is affected.

In addition to Marvis Actions, Juniper Data Center Assurance also provides Marvis Conversational Interface (CI), which supports natural language processing (NLP), to search documentation in Juniper Networks Documentation and Knowledge Base. Network administrators can use Marvis CI to ask queries in natural language and Marvis CI will look up relevant documentation repositories and generate the answer. This enables network administrators to resolve the issues faster without having to manually navigate through different documents to obtain the required information.

About the Marvis Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 114

To access this page, click **Marvis**.

Tasks You Can Perform

- View data center events—View total number of events recorded. You can click an event type to view information about the events reported for the selected event type. You can view more information about the event and the recommended resolution.
- Cross-launch Apstra Data Center Director—Click **View More** from the event details to view more information about the event. From the pop-up that appears, click **Details** to launch Apstra Data Center Director and troubleshoot the event.
- View summary of the events reported for the last seven days.
- Search for relevant information—Click **Ask a question** and enter your query. Marvis searches for the requested information in documentation and populates the results.

View Data Center Events in Marvis Actions

The Marvis Actions dashboard in the Juniper Data Center Assurance application provides network administrators a high-level view of the data center events received from the data center managed by Apstra Data Center Director.

1. Log into Juniper Data Center Assurance.


The Marvis Actions dashboard appears displaying the total number of events under each predefined event category.

2. Click an event category to view all the events recorded for that event category.

3. Click an event to view details of that event.

Marvis Actions dashboard displays the event details along with the recommended action to resolve the event as shown in [Figure 1 on page 115](#).

Figure 6: Marvis Actions Dashboard



4. Click **View More**.

More details about the event such as the site where the event is reported, the device affected, and the date and the time of the event are displayed.

You can mark the status of the event as **In Progress** to indicate that an administrator is working on resolving the event. The event information is removed automatically after the issue is resolved.



NOTE: To view the most recent data center events in Marvis Actions, you must refresh the page by using the **Refresh** button in the browser window. Marvis Actions dashboard does not refresh automatically.

Event Types Displayed in Marvis Actions

Marvis Actions in Juniper Data Center Assurance displays data center events received from a data center managed by Apstra Data Center Director. It also displays the alarms received from a vCenter and provides information about the health and connectivity of the virtual infrastructure in the data center network.



NOTE: Though the intent-based analytics in Apstra Data Center Director supports several predefined probes, Marvis Actions in Juniper Data Center Assurance supports only a limited set of probes. [Table 35 on page 116](#) lists the predefined probes that Marvis Actions supports.

Marvis Actions doesn't display all the events that occur in the network. It displays only those events that are actionable. There could be events that might need additional troubleshooting to resolve the issue.

Table 35: Events displayed in Marvis Actions Dashboard

Event Type	Events
Layer 1 & 2	• Incorrect Cabling • Bad Optics • Interface Flapping • Link Status Mismatch • Packet Discard
Connectivity	• Missing Routes • BGP Mismatch • LAG Imbalance • MLAG Imbalance • BGP Flapping • EVPN Host Flapping • Type-3 Missing Routes • Type-5 Missing Routes • VXLAN Flow Lists Mismatch

Table 35: Events displayed in Marvis Actions Dashboard *(Continued)*

Event Type	Events
Device	• Config Deviation • Deployment Status Mismatch • Resource Health Issues • Environment Health Issues
Virtual Infra	• Config Mismatch • Missing VLANs • MTU Issues • Non-Redundant Hosts • VM Resource Health Issue • Network Connectivity Lost
Security	802.1x Issues
Traffic Capacity	• Spine Faults • Critical Services Alerts • Hot/Cold Interface Warning • Congestion Notification • Out of Sequence



NOTE: Though the events "Congestion Notification" and "Out of Sequence" are displayed in Marvis Actions, information for these events will be available only in a future release of Apstra Data Center Director.

Marvis AI Assistant Overview

IN THIS SECTION

- [Benefits of Marvis AI Assistant](#) | 119

Marvis AI Assistant (Marvis) is an AI-enabled conversational interface that helps you understand, troubleshoot, and explore Juniper Networks products and services. It combines natural language interaction with specialized AI agents to deliver accurate, contextual, and actionable responses.

Marvis is designed to serve as your first-level assistant. It helps you diagnose issues, understand system and network behavior, and access relevant documentation through a single, intuitive interface.

Marvis uses natural language understanding to interpret your query and then orchestrates the appropriate AI agents to analyze data, retrieve relevant information, and generate responses.

By cross-referencing and collating information from multiple authoritative sources, it enables users to resolve issues faster without having to depend on multiple sources.

When you start a conversation with Marvis, it:

- Uses natural language understanding (NLU) to interpret intent and context
- Simultaneously, Marvis invokes purpose-built AI agents to:
 - Perform diagnostic analysis using available network telemetry, logs, and configuration data
 - Analyze application specific data such as operating system logs, blast radius, crash information
 - Correlate insights across the network to identify patterns, anomalies, or potential root causes
 - Retrieve relevant product documentation and provide contextual guidance
- Maintains conversational persistence across follow-up queries and sessions

The AI agents are invoked in response to a user request and operate within predefined scopes. This enables Marvis to deliver more accurate diagnostics, guided troubleshooting, and provide context-aware recommendations and documentation.

Marvis is designed to assist only with Juniper Networks products and services.

Marvis generates responses supported by citations to authoritative Juniper resources. Users can select citation links to view source materials directly. If a query is unrelated to Juniper Networks products or

services, Marvis informs the user that it is designed exclusively to assist with Juniper-related topics and does not generate speculative or unsupported responses.

To enhance usability and effectiveness, Marvis also provides follow-up questions tailored to the current query, easy access to actions such as copying responses, viewing citations, and resetting conversations.

A sample scenario can look like this; you notice intermittent latency in a data center fabric and ask Marvis to find the root cause.

Marvis interprets your description of the issue using natural language understanding to determine the scope and intent of the query. Based on this understanding, it invokes specialized AI agents to analyze relevant network telemetry, logs, and recent configuration changes. It then correlates anomalies across the fabric and retrieve applicable Juniper documentation. And then presents you guided troubleshooting steps with supporting references and allows you to refine the investigation through follow-up questions within the same conversation.

Benefits of Marvis AI Assistant

- Serves as the first-level assistant to address your queries about Juniper Networks products and services.
- Provides an intuitive user experience by facilitating dialog-based interactions and natural conversational flow.
- Use specialized AI agents to offer guided troubleshooting and deeper analysis and insights.
- Generates contextually aware and dynamic responses based on query and past interactions.
- Allows you to ask follow-up queries and clarifications to refine your initial query and its response.
- Provides citations to authoritative resources based on which the response is generated.

Find Information and Troubleshoot Juniper Products and Services Using Marvis AI Assistant

Marvis AI Assistant (Marvis) is a conversational AI assistant that enables you to ask questions in natural language and receive quick, contextual, and actionable responses about Juniper Networks products and services.

Marvis can assist with product information, configuration guidance, troubleshooting, and documentation retrieval within a single conversational interface.

For more information, see ["Marvis AI Assistant Overview" on page 118](#).



NOTE: Marvis capabilities are scoped by the application through which it is accessed. Depending on the entry point, Marvis may use different agents, tools, and data sources to generate responses.

To use Marvis to find information and troubleshoot issues:

1. Log into Juniper Data Center Assurance.
2. Click the **Marvis** icon at the lower right corner of the UI.




Marvis AI Assistant chat window appears at the lower right corner of the page.

Click the **Expand** icon in the chat window to open it in full-screen.

3. To start your conversation with Marvis, use the suggested prompt templates available under the Quick Start section.

Figure 7: Marvis Interface



Click the prompt template that best suits your question.

The chatbox is populated with your selected prompt.

Complete the prompt with the relevant details and press **Enter** to get a response.



NOTE: The suggested prompt templates are optional and are provided for convenience.

Once a response is generated, the prompt suggestions disappear.

4. If your question doesn't fit the suggested prompt templates, you can ignore them and enter your question directly in the chatbox.

Enter the search string or keyword that you want to look up. For example, "Juniper Data Center Assurance supported devices".

You can also enter questions and statements just like you would in a normal conversation. In addition, you can ask follow-up questions and clarifications to refine the query and response.

For example, you can ask a general question like, "How to configure BGP?" and follow it up with a more specific question like, "How to do it on MX204?".

Based on your input, Marvis interprets the intent of the query and applies specialized AI agents to analyze relevant data sources, retrieve applicable documentation, and generate a contextual response. Depending on how the query is processed, the AI agent or tool used to generate the response is displayed at the top of the response.

You can also paste error messages or logs into the chatbox for assistance. Marvis uses agent-driven diagnostics to provide explanations, troubleshooting steps, and links to relevant documentation to help resolve the issue.

Marvis also allows you to copy the query response, view citations, and reset the conversation.

An action toolbar appears at the end of each response, providing quick access to the following functions:

- Copy to Clipboard—Click the copy icon to copy the content of your query, the generated response, and related citations to the clipboard.
- View Citations—Click the citations icon to view the list of sources on which the generated response is based. Click the citation links to explore the source materials.
- Reset Conversation—Click the reset conversation icon to start a new conversation.



NOTE: Resetting the conversation clears your conversation history and starts a new session.

5. You can close Marvis window if your query is resolved, or continue the conversation if you have any follow-up queries.

Marvis suggests related questions that can help you refine your query and explore additional information.

Click a question from the Related Questions section to populate it in the chat and continue the conversation or enter your follow-up question in the chatbox.

6. To start a new conversation, click **New Thread**. Alternatively, press **Ctrl + K**.

You can view and manage your conversation history from the sidebar.

The **Threads** section displays your conversation history, grouped by day.

Click **Delete** to remove conversations from history. Alternatively, click the three dots on a specific thread and select **Delete**.

12

CHAPTER

Monitor and Troubleshoot Data Center Events

IN THIS CHAPTER

- [Adopt Juniper Apstra Edge | 124](#)
 - [Verify DC Assurance Setup | 128](#)
 - [Upgrade Apstra Edge from Juniper Data Center Assurance | 134](#)
 - [About the Dashboard Page | 136](#)
 - [Data Center Network Topologies on DC Assurance | 148](#)
 - [Access Apstra Data Center Director from Juniper Data Center Assurance | 158](#)
 - [Monitor and Troubleshoot Data Center Events from Mist | 160](#)
-

Adopt Juniper Apstra Edge

IN THIS SECTION

- [Apstra Edge Overview | 124](#)
- [Adopt Apstra Edge in Juniper Data Center Assurance | 124](#)

Apstra Edge Overview

Apstra Edge is a hardware-agnostic virtual device that runs within a container in the data center. Apstra Edge functions as a proxy device in a data center managed by Apstra Data Center Director. It receives event and anomaly information from Apstra Data Center Director and forwards that data to Juniper Data Center Assurance.

Apstra Edge is a required component for enabling DC Assurance and its supported features, like [Application Awareness](#), [Impact Analysis](#), and [Service Level Expectations](#).

To enable Apstra Edge to send event information and flow data to DC Assurance, you need to adopt Apstra Edge in DC Assurance.

Adopt Apstra Edge in Juniper Data Center Assurance

Data Center Assurance connects to Apstra Data Center Director through Apstra Edge. For, this you need to adopt Apstra Edge in DC Assurance. After Apstra Edge is adopted successfully, DC Assurance generates a registration code for the edge. This *registration code can be used only once* and must be configured in Apstra Data Center Director to complete edge registration.



CAUTION: The edge registration code is for one-time use only. After the Apstra Edge registers with DC Assurance on first boot, the code is no longer valid. After the first reboot, Apstra uses a stored secret.

The docker container on which the edge runs can be installed anywhere in the data center - either on the same virtual machine (VM) on which Apstra is running or on a different VM. If edge is installed on the same VM as Apstra is installed, you need to take care during upgrades to ensure that the edge installation is not removed.

After the edge is initialized and configured, the status of the edge will be displayed in Juniper Data Center Assurance as **Registered**. The **Cloud Connectivity** status indicates that the edge is connected with the cloud and **Apstra Connectivity** status indicates that the edge is connected with Apstra Controller. That is, you can monitor the status of connectivity between Juniper Data Center Assurance and Apstra Edge, and between Apstra Edge and Apstra from Juniper Data Center Assurance UI without having to access Apstra.

To adopt Apstra Edge:

1. Log in to Juniper Data Center Assurance.
2. Navigate to the **DC Edges** page.
3. Click **Adopt Apstra Edge**.

The **Adopt Apstra Edge** dialog box opens.

Figure 8: Adopt Apstra Edge

Adopt Apstra Edge [X]

Edge Name *

Management URL *

Username *

Password *

Flow Server Details ☒

Flow Server Name *

Management URL *

Username *

Password *

+

vCenter Details ☒

vCenter Name *

Management URL *

Username *

Password *

+

Apstra Edge Download location

4. Enter information about the Apstra Edge device.

Table 36: Apstra Edge Device Details

Field	Description
Edge Name	Unique name to identify Apstra Edge configured in the data center.
Management URL	URL to connect to the Apstra instance. For example, https://10.28.52.3
Username	Username to log into the Apstra instance.
Password	Password to log into the Apstra instance.

5. Enable the **Flow Server Details** toggle button and enter the Flow Server information.

Enabling this option allows DC Assurance to collect flow data from the Apstra instance. After the details for **Flow Server Details** are configured, you can view information about the data center network topology, hosts and endpoints, services, and anomalies, if any, on the **App Aware** and the **Impact Analysis** pages.



NOTE: If you have multiple Apstra Flow servers that collect flow data from their respective sites configured to a single Apstra Edge device, you can add each flow server to the Apstra Edge device and receive the site flow data on Juniper Data Center Assurance.

After the Apstra Edge device is registered and connected, select the Edge device, click **More** on the top right corner, and select **Modify Edge**. Click on the "+" icon under the Flow Server section and enter details for the additional flow servers. You can configure as many flow servers as required.

Table 37: Flow Server Configuration

Field	Description
Flow Server Name	Name of the Apstra Flow server.

Table 37: Flow Server Configuration (*Continued*)

Field	Description
Management URL	OpenSearch REST API URL to connect to the Apstra Flow server. The URL must be in the format <code>https://ip-address:port-number</code> . For example, <code>https://10.28.52.4:9200</code> .
Username	Username to log into Apstra Flow server.
Password	Password to log into Apstra Flow server.



NOTE: The Apstra Flow information collected from the Apstra instance is stored in Juniper Data Center Assurance database for 30 days.

6. Enable the **vCenter Details** toggle button and enter the vCenter information.

Enabling this option allows DC Assurance to collect information from the vCenter about the virtualized environments and the alarms detected on VMs. After the details for vCenter are configured, you can view information about the VM layer on the **App Aware**, the **Impact Analysis**, and the **Dashboard Topology** pages. You can configure parameters for multiple vCenters.

Table 38: vCenter Configuration

Field	Description
vCenter Name	A unique, user-defined name of the vCenter.
Management URL	REST API URL to connect to the vCenter server. The URL must be in the format <code>https://ip-address:port-number</code> .
Username	Username to log into vCenter server.
Password	Password to log into vCenter server.



NOTE: The **Apstra Edge Download Location** field displays the location for downloading Apstra Edge software. Visit <https://support.juniper.net/support/downloads/?p=apstra> to download Apstra Edge software.

7. Click **Adopt**.

The newly added edge is listed on the **Inventory** page with the status **Unregistered** and **Disconnected**.

8. Install and set up Apstra Edge.



NOTE: If you are using Apstra Data Center Director 5.1.0 or later, you can set up Apstra Edge from within the Apstra Data Center Director GUI. For more information, see [Juniper Data Center Assurance Onboarding Guide](#).

If you are using Juniper Apstra 5.0.0 or earlier, follow the steps provided in [Juniper Apstra Edge Setup Guide](#).

The status of Apstra Edge will be displayed as **Registered** and **Connected**, indicating that the edge is registered and connected. A site corresponding to the blueprint deployed in the data center is created on the Sites page.



NOTE: Deleting the adopted edge from Juniper Data Center Assurance also removes the connectivity with edge. Once the edge is deleted, Juniper Data Center Assurance will not be able to receive event information from Apstra. To continue receiving events, you need to perform the registration process again.

Verify DC Assurance Setup

IN THIS SECTION

- [Juniper Data Center Assurance Setup Validation Tool | 129](#)
- [Pre-Flight Checks | 129](#)

The Apstra Edge device forwards event and anomaly information from the data center to Juniper Data Center Assurance. To ensure that DC Assurance receives insights and anomaly data from the data center, relevant configurations must be made in Apstra and on the Edge. You can verify these configurations by using:

- **Juniper Data Center Assurance Setup Validation Tool**—verifies the data center fabric connectivity and the device configurations.
- **Pre-flight Checks**—verify the Edge version and blueprint configurations on Apstra.

Juniper Data Center Assurance Setup Validation Tool

After the Apstra Edge device has been installed and set up, you can run the Juniper Data Center Assurance Setup Validation Tool. This tool performs a series of tests to verify that you have configured the various devices and services of your DC fabric correctly.

For information about the Juniper Data Center Assurance Setup Validation Tool, see [Download and Run the Juniper Data Center Assurance Setup Validation Tool](#).

Pre-Flight Checks

The pre-flight checks feature verifies whether certain prerequisites and configuration requirements are met and provides actionable guidance on how to resolve potential issues. Executing these checks enables you to validate the environment, identify issues that might affect the performance of DC Assurance features, and resolve them.

The pre-flight checks feature helps you verify whether:

- the blueprints are configured properly.
- the probes that relay data center events and insights from Apstra Data Center Director to DC Assurance are enabled.
- the Edge device is running the minimum software version required to support the latest DC Assurance features.
- the Edge device is connected to the Apstra Flow Server.
- the system times on the Apstra Flow Server and Juniper Data Center Assurance are in sync.
- the stream receivers in Apstra are present and connected.

You can run pre-flight checks at the Edge level or for specific blueprints. If you run the pre-flight checks at the blueprint level, tests are run only on the selected blueprints. If you run the pre-flight checks at the Edge level, tests are run on the Edge as well as all the blueprints enabled on the selected device.

The pre-flight checks feature runs the following tests:

Table 39: Pre-flight tests

Component	Pre-Flight Test	Description
Apstra Edge	Edge Version Compatibility Check	Checks whether the onboarded Apstra Edge is running the minimum required version to support DC Assurance features.
	Stream Receivers Creation Check	Checks whether the following stream receivers are present and connected: • Alerts • Perfmon
	Apstra Flow Server Connection Check	Checks whether the Apstra Edge device and the Apstra Flow Server are connected.
	Time Stamp Check between Apstra Flow Server and Juniper Data Center Assurance	Checks whether the timestamps on the Apstra Flow Server and Juniper Data Center Assurance are in sync.
Blueprint	Stream CPU and Memory Utilization Probe	Checks whether streaming is enabled for disk utilization, CPU utilization, and memory utilization data.

Table 39: Pre-flight tests (Continued)

Component	Pre-Flight Test	Description
	SLE Probes	Checks whether the following probes are enabled on Apstra: • Device Environmental Checks • Drain Traffic Anomaly • Interface Flapping (Fabric Interfaces) NOTE: You can enable any of the following optional probes to receive more insights from Apstra: • BGP Monitoring • ECMP Imbalance (Fabric Interfaces) • EVPN Host Flapping • EVPN VXLAN Type-3 Route Validation • EVPN VXLAN Type-5 Route Validation • Hot/Cold Interface Counters (Fabric Interfaces) • Hot/Cold Interface Counters (Specific Interfaces) • Interface Flapping (Specific Interfaces) • Optical Transceivers • VXLAN Flood List Validation When a pre-flight check is run on a blueprint, the SLE Probes section in the Pre-flight

Table 39: Pre-flight tests *(Continued)*

Component	Pre-Flight Test	Description
		Information page displays the list of optional probes that have not been enabled on Apstra.

Before you perform a pre-flight check, ensure that:


- the Apstra Edge is registered.
- the Apstra Edge is connected to the cloud and with the Apstra Controller.
- any upgrade process on the Apstra Edge is complete.

To run a pre-flight check:

1. Log in to Juniper Data Center Assurance.
2. Navigate to the **DC Edges** page.


The onboarded Edges are listed along with their connectivity status, the current version of software running on the Edge, available versions for upgrade, upgrade status, and results from the previous pre-flight check.

Figure 9: Onboarded Apstra Edge Devices



3. Select the Edge or the blueprint that you want to verify.

Figure 10: Pre-flight Check Selection



4. Click **Run Pre-flight Checks** on the top right of the table.

You can see the status of the check in the **Pre-flight Results** column.


5. Click the link in the **Pre-flight Results** column to view the results. The **Pre-flight Information** page opens.

This page displays information about whether the device or blueprint passed or failed the tests. Select the failed test to view more information and guidance on how to resolve the issue.



NOTE: The **SLE Probes** section in the **Pre-flight Information** page for a blueprint also displays a list of optional probes that you can enable to receive more detailed insights from Apstra.

Figure 11: Pre-flight Information





NOTE: You can perform pre-flight checks on an Apstra Edge device or on blueprints at any time after the device has been onboarded.

Upgrade Apstra Edge from Juniper Data Center Assurance

SUMMARY

This document describes how to upgrade Apstra Edge remotely from Juniper Data Center Assurance.

To access the advanced features in Juniper Data Center Assurance and to obtain accurate event and anomaly information, you must upgrade Apstra Edge to a version that supports those features. You can upgrade Apstra Edge manually or from the Juniper Data Center Assurance UI. This topic describes how to upgrade Apstra Edge remotely from Juniper Data Center Assurance. You can view the Apstra Edges that require upgrade and the available software images in the **DC Edges** page.



NOTE: You must be a superuser to upgrade Apstra Edge.

Before you begin the remote upgrade process, please ensure that:

- the Apstra Edge container is deployed.
- the Apstra Edge is connected to the cloud and with the Apstra Controller.

To upgrade Apstra Edge:

1. Log in to Juniper Data Center Assurance.
2. Navigate to the **DC Edges** page.

The Edges are listed along with their connectivity status, the current version of software running on the Edge, available versions for upgrade, and upgrade status.



NOTE: Ensure that Apstra Edge is connected to the cloud to enable remote upgrades when a new version is available.

3. Select the check box for the Edge(s) that you want to upgrade.

	Name	Management URL	Registration Status	Cloud Connectivity	Apstra Connectivity	Flow Server Status	Current Version	Latest Version	Upgrade Status	VCenter Status
☒	edge-1	https://10.82.94.10	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-2	https://10.82.94.11	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-3	https://10.82.94.12	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-4	https://10.82.94.13	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-5	https://10.82.94.14	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	



NOTE: When you select multiple Edges, if the **Upgrade Edge** option is disabled, the tooltip provides information about why one or more of the selected Edges might not be eligible for an upgrade.

4. Click **More** > **Upgrade Edge** on the top right of the table.

	Name	Management URL	Registration Status	Cloud Connectivity	Apstra Connectivity	Flow Server Status	Current Version	Latest Version	Upgrade Status	VCenter Status
☒	edge-1	https://10.82.94.10	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-2	https://10.82.94.11	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-3	https://10.82.94.12	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-4	https://10.82.94.13	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	
☒	edge-5	https://10.82.94.14	Registered	Connected	Connected		0.13.3	0.13.4	New Version Available	

The **Upgrade Apstra Edge** dialog box opens with the list of selected Edges and the latest available versions for each Edge.

5. Click **Upgrade**.

A confirmation window appears.

6. Click **Yes** to begin the upgrade process.

You can track the upgrade progress in the **Upgrade Status** column. Hover over the progress bar to view the current state of the upgrade process. The status of the upgrade is automatically updated every 30 seconds. You can also click the refresh icon on the top right of the table to view the latest status.

Once the upgrade process is complete, the **Upgrade Status** column shows the status as **Complete** and the **Current Version** column displays the version of software newly installed on the Edge.

After an upgrade, the Edge is monitored for the first 20 minutes. If the newly installed version fails during this period, the software version is reverted to the version that was previously running on the Edge. Beyond this time period, if the Edge crashes, the docker container restarts with the newly installed software version.

If the upgrade process fails for an Edge, the **Upgrade Status** column shows the status as **Failed**.

In both failure scenarios, you can retry upgrading the Edge by selecting the Edge and performing the upgrade process again.

About the Dashboard Page

IN THIS SECTION

- [Overview Tab | 136](#)
- [Widgets Displayed for Site Groups and All Sites | 137](#)
- [Widgets Displayed for a Single Site | 139](#)
- [Anomalies | 142](#)
- [Geographic View Tab | 142](#)
- [Topology Tab | 144](#)
- [Accessibility on the Dashboard | 146](#)

The Juniper Data Center Assurance dashboard provides operational insights about the performance of the data center network and the hosted services. These insights enable network administrators to quickly respond to network issues that impact end user experience.

To access this page, click **Dashboard** from the main menu. The dashboard provides information in the three tabs: **Overview**, **Geographic View** and **Topology**.

Overview Tab

The **Overview** tab provides three views of the data center network: individual site, site group, and for all the sites in the organization that you have access to based on your user role. The site-level view provides information about the devices in the site, their role, anomalies in the devices, impact analysis information, and so on, displayed in a set of predefined widgets. The information in these widgets gets updated on the basis of the information available for the selected site.

You can select the widgets that you want and customize how they are displayed in the **Overview** tab. Click the **Edit** (pencil) icon at the top right of the page to show or hide widgets from the **Overview** tab. Also, you can resize or reorder the widget according to your layout preferences. Click the ellipsis (...) at

the top right of any widget to refresh the displayed information, resize, or to reorder the layout of the widget as shown in [Figure 12 on page 137](#).

You can also create a custom dashboard to monitor specific insights according to your network's requirements, by clicking **Add Tab**. You can then add the required widgets.



Video: [Dashboard](#)

Widgets Displayed for Site Groups and All Sites

In the **Overview** tab, you can also view information for a site group or for all the sites in the organization that you have access to based on your user role. Use the filter at the top of the page to select a site group or **All Sites**. When a site group or **All Sites** is selected, the widgets display information about Apstra Controllers, version of Apstra installed, and so on. [Figure 12 on page 137](#) displays how the **Overview** tab looks when a site group or All Sites is selected.

Figure 12: Overview Tab for All Sites




Table 40: Widgets Displayed for Site Group and All Sites

Widget Name	Description
Apstra Cluster Health	Displays the total number of controllers within an Apstra cluster in the entire organization or a site group. It also provides a summary of the CPU, memory, and disk utilization, site count, worker count, and remote server count for each controller. Select an individual controller to expand the view and see details of its worker nodes and flow servers. You can also select a worker node or a flow server to view its historical data.
Version Vs Instance	Number of Apstra instances running and their versions.
Top 3 Devices by Port Count	Displays a graph listing the top three devices by number of ports.
Summary Inventory	Displays hierarchical information about the superspines, spines, leaf, and host devices in the organization or site group.
Inventory	Displays the total count of devices in the data center and their roles.
Anomalies	Displays the total count of anomalies categorized by their severity levels, such as critical, warning, and low. The tile also lists the top three critical anomalies that need immediate attention.
Port Activity	Displays the total number of ports in the organization or site group along with a graph of used and unused ports.
Summary Overview	Provides information about the total number of sites, devices, services, and the total number of anomalies generated.


Table 40: Widgets Displayed for Site Group and All Sites *(Continued)*

Widget Name	Description
Potentially Impacted Clients	Provides the total number of clients that are impacted because of an anomaly and the total number of services affected.

Widgets Displayed for a Single Site

You can select a single site to view site-level information in the **Overview** tab. The widgets on the page provide information specific to the site you selected. [Figure 13 on page 140](#) shows the widgets displayed on the Overview tab for a single site.

Figure 13: Widgets Displayed for a Single Site



In [Figure 13 on page 140](#), the **Key Anomalies** widget and the **Anomalies** tab in the right pane do not show any data because no anomalies are reported in the site you selected. The widgets get updated depending on the information available for the selected site. [Table 41 on page 141](#) describes the widgets displayed for a single site.

Table 41: Widgets Displayed for a Single Site

Widget Name	Description
Apstra Cluster Health	Displays the controllers within an Apstra cluster for the selected site with all worker nodes and remote servers associated with it. It also provides a summary of the CPU, memory, and disk utilization, site count, and worker count for the controller and each worker node and flow server. You can select a worker node or a flow server to view its historical data.
Top 3 Devices by Port Count	Displays a graph listing the top three devices by number of ports.
Top 3 Services By Volume	Displays a graph listing the top three services by volume of data.
Summary Overview	Displays the number of devices and active services in the site and the count of anomalies generated.
Summary Inventory	Displays hierarchical information about the superspines, spines, leaf, and host devices in the site.
Inventory	Displays the total number of devices in the data center and their roles.
Anomalies	Displays the top three key anomalies along with the total count of anomalies generated in the site.
Port Activity	Displays the total number of ports in the site along with a graph of used and unused ports.
Site Details	Displays the Apstra version, number of physical and logical systems, and the anomalies reported.
Potentially Impacted Clients	Displays the total number of clients and services impacted.

Table 41: Widgets Displayed for a Single Site *(Continued)*

Widget Name	Description
Network Elements (Non-switch)	Displays the number of different non-switch elements in the network. Also shows the number of elements with anomalies and those operating normally in each element type.


Anomalies

The **Anomalies** tab in the right pane displays the anomalies that have occurred and also indicate how the anomalies are related to a key anomaly. An anomaly might trigger multiple related events. Expand the key anomaly to view the detailed list of anomalies related to the key anomaly. Often, the related events that are generated get resolved automatically when the event that triggered these anomalies is resolved. Such an event or anomaly is termed a key anomaly. Identifying and flagging key anomalies help administrators resolve issues faster. Identifying and flagging key anomalies also help reduce alert fatigue, which is a condition that administrators might face due to overwhelming number of anomalies of different levels of severity. Alert fatigue might lead to delayed response to anomalies or sometimes ignoring critical anomalies.

Geographic View Tab


The **Geographic View** tab helps view the location of the sites on the map, number of devices, and the anomalies if any in the devices. The site icon on the map appears in blue if the site is operating normally. The icon appears as red if anomalies are reported in the site as shown in [Figure 14 on page 143](#). Hover over the icon to view the site details. In the pop-up that appears, you can view site information, such as the number of devices in the site, the version of Apstra installed in the site, and number of Apstra instances.

Figure 14: Geographic View Showing Anomalies




Click the icon to zoom in to the exact location of the site on the city map. Click **View Impact Analysis** as shown in [Figure 15 on page 144](#). You are prompted to select the site. After you select the site, you are redirected to the **Impact Analysis** page.

Figure 15: Dashboard - Geographic View




Topology Tab

The **Topology View** tab provides a graphical representation of the physical topology of the data center fabric, providing information about how devices are connected to one another, their roles, and the count of anomalies, if any, generated in the devices. The topology view displays the data center network in three layers - host, leaf, and spine. You can also view the VM connections in the topology by enabling the **Show VM layer** toggle button.

Figure 16: Topology View

In addition to this, you can view details of specific nodes in the topology. Click a node to view device information such as CPU and memory utilization in the right pane as shown in [Figure 17 on page 146](#).

Figure 17: Topology View - Device Information



Accessibility on the Dashboard

The DC Assurance dashboard is compliant with the [Web Content Accessibility Guidelines 2.1 \(WCAG\)](#) from the [World Wide Web Consortium \(W3C\)](#).

W3C is an organization that helps develop Web standards and guidelines to accommodate the growing diversity of people, hardware, and software. WCAG is a set of guidelines developed with a goal of providing a single shared standard for web content accessibility that meets the needs of individuals, organizations, and governments internationally. These guidelines explain how to make web content more accessible to a wide range of audiences, including users viewing with assistive tools.


The dashboard is accessible to users who employ screen readers such as [Non-Visual Desktop Access \(NVDA\)](#) for devices running Windows or VoiceOver that is built into Apple devices. You can navigate and interact with dashboard elements using a screen reader in combination with the keyboard shortcuts listed in [Table 42 on page 147](#).

Table 42: Keyboard Shortcuts

Action	Shortcut Key
Forward Navigation	Tab
Backward Navigation	Shift + Tab
Select, Deselect, Open	Enter / Space
Navigate between options	ArrowUp / ArrowDown
Close	Esc

On using the keyboard navigation, the selected element is highlighted as shown in [Figure 18 on page 148](#).

Figure 18: Keyboard Navigation



Data Center Network Topologies on DC Assurance

IN THIS SECTION

- [3-Stage Clos Architecture | 153](#)
- [3-Stage Clos Architecture with Virtual Devices | 153](#)
- [Collapsed Data Center Fabric | 154](#)
- [5-Stage Clos Architecture | 155](#)
- [Network Architecture with Access Layer Support | 156](#)

A data center network consists of various components like servers and storage devices. Data center networking architecture refers to the design and organization of the network systems within a data center. A data center network must be designed to strike a balance between reliability, performance, agility, scalability and cost.

You can use the templates in Apstra Data Center Director to create a blueprint and design your data center network. For more information, see [Create Blueprint](#) in the *Apstra Data Center Director User Guide*.

Juniper Data Center Assurance displays the data center network topologies on the **Topology** tab on the **Dashboard** page, the **Service Topology View**, and the **Impact Analysis** pages. DC Assurance also provides actionable insights about the device health and performance, events, and impacts of those events in a data center network.

[Figure 1 on page 150](#) shows the **Overview** tab on the **Dashboard** page. The various widgets provide information about the devices in the site(s), their role, anomalies in the network and so on.

Figure 19: Overview tab on the Dashboard page





Figure 2 on page 151 shows the **Impact Analysis** page with information about data center events and anomalies along with the impact of these anomalies on the data center network.

Figure 20: Impact Analysis page



The network administrators use these insights to monitor the network and resolve issues to ensure that services are not impacted. Use the filters on the **Service Topology View** and **Impact Analysis** pages to select the layers to be displayed.

Figure 21: Show Layers Filter

Time Increment

15 min

Time Range

Nov 24, 2025 1:00 PM - 1:15 PM

Show Layers

External Router x

Service x VM x

Host x Access x

Leaf x Spine x

Super Spine x

- ✓ Service
- ✓ VM
- ✓ Host
- ✓ Access
- ✓ Leaf
- ✓ Spine
- ✓ Super Spine
- ✓ External Router

Currently, DC Assurance supports the following data center network architectures:


- ["3-Stage Clos Architecture" on page 153](#)
- ["3-Stage Clos Architecture with Virtual Devices" on page 153](#)
- ["Collapsed Data Center Fabric " on page 154](#)
- ["5-Stage Clos Architecture" on page 155](#)
- ["Network Architecture with Access Layer Support" on page 156](#)
- ["Network Architecture with External Router Support" on page 157](#)

3-Stage Clos Architecture

The 3-stage Clos Architecture provides organizations with a data center network that is fast, adaptable to change, scalable, and reliable. This topology consists of 3 components - server leaf switches, border leaf switches, and spine switches.

DC Assurance displays information and provides insights about the spines, leaf devices, hosts, and services in the data center network.

Figure 22: 3-Stage Clos Topology




3-Stage Clos Architecture with Virtual Devices

The 3-Stage Clos Architecture with Virtual Devices combines physical networks with network virtualization. This topology builds on a 3-stage fabric as the foundation for virtual network integration. In this design, the host layer is connected to different services through VMs.

DC Assurance collects information about the VMs in the data center and displays the VM layers in the topology views. To enable this feature, you must configure vCenter server on the Adopt Apstra Edge page. You can configure multiple vCenters.


Figure 23: 3-Stage Clos Topology with Virtual Devices



Collapsed Data Center Fabric

The collapsed data center fabric is a two-switch network fabric architecture designed for small network deployments. In collapsed fabric, switches perform the roles of spine, leaf, and border leaf switches. This design allows for high availability network deployments with minimum switch hardware. However, the hardware constraints limit the scalability of this architecture.

Figure 24: Collapsed Data Center Fabric Topology



5-Stage Clos Architecture


The 5-stage Clos architecture is similar to the 3-stage data center design with the exception of a superspine layer. This design allows for scaling large scale data center design with requirements for large datastores and the compute nodes that need to connect to the data storage. In this topology, an additional layer of superspine devices interconnects multiple pods consisting of spine-and-leaf layers (similar to 3-Stage fabric), thus creating a large network.

DC Assurance displays information and provides insights about the superspine layer, spines, leaf devices, hosts, and services in the data center network.



NOTE: Juniper Data Center Assurance currently does not support the display of pods.

Figure 25: 5-Stage Clos Topology




NOTE: If your 5-Stage Clos fabric is rendered as a 3-Stage Clos fabric on DC Assurance, you must make a minor change in the site blueprint (for example, change the blueprint description) in Juniper Apstra Edge device. The data center network is rendered as a 5-Stage Clos fabric displaying the superspine layer after DC Assurance synchronizes data with the Edge device.

Network Architecture with Access Layer Support

Access switches connect end-user devices to the data center network. These switches also manage and direct data traffic to ensure that the packets are sent to the correct destination. Access layer switches make the data center networks more scalable, provide security to the network devices, and reduce operational complexity of the network by simplifying device connections.

Apstra supports access switch layer in 3-stage Clos, 5-stage Clos, and collapsed fabric network architecture. DC Assurance displays information about the health and status of the access switches in the network along with the other network components, thus helping prevent potential issues and manage network infrastructure.

Figure 26: Network Architecture with Access Switches




NOTE: DC Assurance displays information only about Juniper devices on the access layer.

Network Architecture with External Router Support

Apstra allows you to connect your data center network to external systems that are not managed by Apstra. These systems are called External Generic Systems. These systems provide connectivity between two or more data centers as Data Center Interconnect (DCI) connections to share data and resources. The external systems can be firewalls, external routers, and so on. The external generic systems are not part of a rack topology and are typically connected to a border leaf. For more information, see [Create External Generic System](#) in the *Apstra Data Center Director User Guide*.

DC Assurance supports visualization of external routers connected to a border leaf of the data center network architecture and provides information about the health and status of the external router. It displays the data center clients for which the external router acts as a gateway and also displays flow data information between the data center and the external router. This information helps network administrators design data center networks and make informed decisions to optimize network operations.

Figure 27: Network Architecture with External Routers



NOTE: Since the external systems are not Apstra-managed, DC Assurance does not display metrics or anomaly data for the external routers.

RELATED DOCUMENTATION

[Application Awareness Overview](#)

[Impact Analysis Overview](#)

Access Apstra Data Center Director from Juniper Data Center Assurance

Make sure that you have completed the following:

- Installed and set up Apstra Edge.

- Registered Apstra Edge in Juniper Data Center Assurance.

Juniper Data Center Assurance provides detailed insights into the operations of the data center. The events that Juniper Data Center Assurance receives from Apstra Edge are organized into six event types and displayed in the Marvis Actions dashboard. Each leg in Marvis Actions dashboard corresponds to an event type. You can click the Marvis Actions leg for an event type to view all the events for the selected event type. Click an event to know the reason for the event and recommended resolution. You can then launch Apstra Data Center Director from Juniper Data Center Assurance to troubleshoot and resolve the event. As the information in the Marvis Actions dashboard is real time, network administrators can resolve events even before they impact network traffic.

1. Log into Juniper Data Center Assurance.

Marvis Actions dashboard is displayed. The Marvis Actions dashboard for data center displays the events categorized by event type. Each Marvis Actions leg displays the total number of events reported for an event type.

2. Click the Marvis Actions leg to view the events for the selected event type.

The various events reported for the event type are displayed.

3. Click an event.

More details about the event such as the site where the event is reported, the device affected, and the date and the time of the event are displayed.

4. Click **View More** to view more information about the event.

A pop-up appears displaying more information about the event.

The screenshot displays the Juniper Apstra MARVIS ACTIONS dashboard. The main view shows a central 'ACTIONS' hub with 148 events, branching into categories like '32 Layer 1 & 2', '33 Connectivity', '12 Traffic Capacity', and '0 Security'. Under '33 Connectivity', there are sub-categories: 'Incorrect Cabling' (8), 'Bad Optics' (22), 'Interface Flapping' (0), and 'Link Status Mismatch' (2). A pop-up window titled 'Incorrect Cabling Details' is open, showing event information for 'spine_leaf3'.

Incorrect Cabling Details Pop-up:

- Start Time: Apr 16, 2024, 09:34 PM
- Hostname: leaf3
- Local ifName: ge-0/0/1
- Role: spine_leaf
- Blueprint label: Great Plains DC-1
- Expected hostname: spine2
- Expected ifName: ge-0/0/1
- Alert id: 58ef081d-25ca-411c-a215-8c1fc119d79
- Remote ifName: ge-0/0/2
- Received time: 1713283463584
- Local hostname: leaf3
- Remote hostname: spine2
- Alert type: CABLE_PEER_MISMATCH_ALERT

INCORRECT CABLING Table:

Site	Device	Details	Date	Action
☐ Great Plains DC-1	leaf3	Incorrect Cabling Detected View More	Apr 16, 2024 9:34 PM	Open
☐ Great Plains DC-1	spine2	Incorrect Cabling Detected View More	Apr 16, 2024 9:34 PM	Open
☐ Great Plains DC-1	spine2	Incorrect Cabling Detected View More	Apr 16, 2024 9:34 PM	Open
☐ Great Plains DC-1	leaf2	Incorrect Cabling Detected View More	Apr 16, 2024 9:34 PM	Open

5. Click **Details** to access Apstra Data Center Director application.

Apstra Data Center Director opens in a new browser window or tab.

Monitor and Troubleshoot Data Center Events from Mist

If you manage your enterprise network using Mist and data center using Apstra, you can monitor data center events too from Mist by linking the organization in Mist with the organization in Juniper Data Center Assurance. Once the organization in Mist is linked with the organization in Juniper Data Center Assurance, you can view the total number of data center events in **Data Center/Application** category in Marvis Actions in Mist. You can also access Juniper Data Center Assurance by clicking the **Data Center Actions** event type under **Data Center/Application** and view more detailed information about an event.



NOTE: You must be a user with the superuser role to link the organizations in Juniper Data Center Assurance and Mist.

Make sure you have the following:

- Login credentials to access Juniper Data Center Assurance
- Login credentials to access Mist
- API token generated in Juniper Data Center Assurance

To access Juniper Data Center Assurance from Mist:

1. Log in to Juniper Data Center Assurance and navigate to the **Organization > Settings** page.
2. Locate the API Token tile and click **Create Token**.

API Token

Create Token

Name	Access	Site Access
------	--------	-------------

3. Enter a name for the token and click **Generate**.

Create Token ✕

Please save your key to a safe place. You will see the key only once upon creation. You won't be able to retrieve it later

Name

Access Level

☒ **Super User**
Full access to organization and all its sites, able to create new sites, and unable to manage API tokens

☐ **Network Admin**
Full access to selected sites

☐ **Observer**
Monitor only access to selected sites

☐ **Helpdesk**
Helpdesk monitoring and workflow for selected sites

Site Access

All SitesSite GroupsSpecific Sites

Key

📋

DoneCancel

- After the token is generated, click the **Copy** button.



NOTE: Store the API token key in a secure location. You can copy the key only once and cannot be retrieved later.

5. Click **Done**.
6. Log in to the Mist portal.
7. Navigate to **Organization > Settings** page.
8. Locate the **Apstra Cloud Services Integration** tile.


The screenshot shows the Juniper Mist APSTRA DEMO interface. On the left is a navigation menu with options like Monitor, Marvis™, Clients, Access Points, Switches, WAN Edges, Mist Edges, Private 5G, Location, Analytics, Site, and Organization. The main content area has a top bar with 'APSTRA DEMO' and a 'Create' button. Below this are several configuration tiles: 'Third Party Token' with a 'Name' field (containing 'X-CP-API-ID') and an 'Add Credentials' button; 'Marvis Minis' with a 'Disable Marvis Minis' checkbox, 'Custom URLs' field, and 'Add Custom URLs' button; 'Application Insights Integration' with a 'Link Account' dropdown and a table with columns 'Status', 'Application', 'Company Name', and 'Limit Daily A'; and the 'Apstra Cloud Services Integration' tile, which is highlighted with a blue border and contains fields for 'Organization ID', 'API Token Name', and 'API Token'.

Enter the following information:

- **Organization ID**—Copy the organization ID from Juniper Data Center Assurance and paste it here.
- **API Token Name**—Enter API token name that you defined in Juniper Data Center Assurance.
- **API Token**—Copy the API token generated in step 3 in Juniper Data Center Assurance and paste it here.


9. Click **Save**.

The organizations in Mist and Juniper Data Center Assurance are linked now. In a few minutes, you'll notice that the **Data Center/Application** event category is active and displays the total number of data center events under the **Data Center Actions** event type.



- Click the **Data Center Actions** event type under **Data Center/Application** to launch Juniper Data Center Assurance.

Juniper Data Center Assurance opens in a new browser window or tab in read-only mode. Log in to DC Assurance to access all features of the application.



13

CHAPTER

Application Awareness

IN THIS CHAPTER

- [Application Awareness Overview | 166](#)
 - [About the App Aware Page | 171](#)
 - [About the Service Topology View Page | 173](#)
-


Application Awareness Overview

IN THIS SECTION

- Benefits of App Aware | 170

The Application Awareness feature in Juniper Data Center Assurance provides a visual representation of the services that are active during a given time period. The visualization also shows how the data center fabric is being utilized by these services. You can select a service and drill down to get detailed information about the traffic flow.

Figure 28: App Aware Page



A data center managed by Apstra Data Center Director is aware of all the services that are running in the data center. The Application Awareness feature provides a visual representation of all the services running in the site, grouped by throughput, volume, or connections for the site you select. The data


displayed on the App Aware page is refreshed every 15 minutes. To view real-time data, click the **On-Demand Data** button on the top right of the page.

You can drill down and view details of specific services, view the network topology, identify any node that has an anomaly, and the hosts connected to the device, which are impacted by the anomaly.



NOTE: You must configure sFlow in the devices in the data center to view flow data in Juniper Data Center Assurance.

Figure 29: App Aware - Topology View




You can select a set of services and examine how nodes are connected to hosts and how those services traverse through the data center fabric. The total count of events and anomalies affecting the nodes is displayed in the topology diagram.



Video: [Application Awareness](#)


You can view a visual representation of the impact of a device or service failure on the clients connected to it. Select a device or service and click **Fail Scenario** on the right pane. The potential failure scenarios are simulated and displayed in the topology.

Figure 30: Service Topology View - Failure Scenario



You can also view how traffic is flowing through the data center network for a specific service. Select a service from the topology and click **Traffic Flow** on the right pane. Once you select the clients you want to monitor from the drop-down, you can see a visualization of application traffic flow.

Figure 31: Service Topology View - Traffic Flow




The **Summary** tab in the right pane provides a summary of the selected service with details such as the number of devices associated with the service and the total number of services.

The search option on the App Aware page enables you to examine the required service without having to drill down from the topology view. You can directly search for the service and view the connections and flow details for that service. The filter option helps you look into the connections between specific leaf and spine nodes and examine the service traffic by selecting specific services and the layers that you want to view.

You can also compare how the services and endpoints performed during two time periods from the **Compare** tab. The flow data collected from the Apstra instance is stored in Juniper Data Center Assurance for 30 days. You can compare the current performance of the services with the last 30 days' flow data. Select **Connection View** to view the connections between the devices and the services associated with the devices. Select **Flow View** to understand how service traffic utilizes the network.

Figure 32: Compare Topology Changes



The flow view on the App Aware page provides information about the flow data between the nodes. A thicker line connecting the devices indicates that more traffic is being exchanged between those devices.

The Application Awareness feature thus provides a complete view of the data center, providing insights about potential issues in the data center network and their impact on the services.

Benefits of App Aware

- Provides a service-to-data center resource mapping, thereby providing visibility about how the services are utilizing the data center fabric. This visibility enables administrators to correlate application issues with potential network issues in the data center devices, helping them to resolve issues faster.
- Provides information about devices or ports that are associated with specific services. This enables administrators to perform impact analysis and pinpoint devices or ports that are causing issues.
- Provides visualization of the network topology displaying how services connect to different endpoints.
- Provides insights about traffic patterns and performance metrics for services and network endpoints.

RELATED DOCUMENTATION

[About the App Aware Page | 171](#)

[About the Service Topology View Page | 173](#)

About the App Aware Page

SUMMARY

The App Aware page provides a bird's-eye view of the devices, services, and traffic flow in the data center network. The Application Awareness feature enables you to monitor how the devices and the services are performing and helps you resolve issues before they impact end user experience.

IN THIS SECTION

- [Tasks You Can Perform | 171](#)

To access this page, click **App Aware** from the main menu.

The App Aware page displays services grouped by throughput, volume of data utilized, or by connections. You can use the filter to change the view.

Tasks You Can Perform

- View real-time information about how traffic is handled by the devices in the data center. Click the **On-Demand Data** button to see latest data. The App Aware page is automatically refreshed every 15 minutes.
- View the topology of the path that a specific service utilizes. Click the **Select** button and select one or more services, and then click **View topology**. The **Connection view** displays how the devices are connected to one another, whereas the **Flow view** displays the traffic flow pattern.
- Search for a specific device or service and view the topology. You can enter the IP address of the device or the port number associated with the service.
- View information about the clients accessing the service. By default, the connected clients are displayed as a group. To see how the clients are connected to the host, click **Ungroup**.
- View a summary of the hosts and clients, their role in the data center topology, and the throughput on the **Summary** pane on the right.

- Filter the information to be displayed on the App Aware page. You can use the filter to display information about specific services and the devices in a site.
- View CPU and memory utilization of the devices. From the Service topology view page, select a device. The CPU and memory utilization graphs and the count of transmitted and received packet at each interface are displayed in the right pane.
- Simulate Fail scenarios—In the Service topology view page, click **Fail Scenarios**. Potential failure scenarios are simulated and displayed in the topology. You can simulate failure scenario at the host level or at the switch level.
- View traffic flow—You can view how traffic is flowing through the data center network. To view traffic flow, select the service and click **Traffic Flow** on the right pane. Once you select the clients you want to monitor, you can see a visualization of application traffic flow.
- Compare the data center topology and the services between two time ranges—Click the **Compare** Tab to identify how many services or clients got added or removed during the selected time range. You can move the slider to view the changes occurred in the network between the two time ranges.

Figure 33: Options on the App Aware Page




Table 43: Options on the App Aware Page

Option	Description
1	On-Demand Data—Click to obtain the latest data from the data center network. The data on the page is refreshed every 15 minutes.

Table 43: Options on the App Aware Page *(Continued)*

Option	Description
2	Search—Search for specific devices or services. Enter the IP address of the device or the port number associated with the service.
3	Grid View/Table view—Change the page view to grid view or table view.
4	Filter—Filter the information displayed on the App Aware page.
5	View topology—View the topology for the selected services. Select one or more services and click View topology .

RELATED DOCUMENTATION

| [Application Awareness Overview](#) | 166

About the Service Topology View Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 174

The Topology service view page displays the topology of the data center network that a service utilizes.

To access this page, click **App Aware** from the main menu and select a select a service.

Tasks You Can Perform

- View the data center network topology that the selected service utilizes.
- View Apstra flow data—You can view traffic patterns between the devices in the data center network.
- View a summary of the topology.
- Simulate failure scenario—Graphical representation of the impact of a device failure on the clients connected to the device. Select a device and click **Fail scenario**. Similarly, you can visualize how a service failure can impact the clients utilizing the service.
- View traffic flow—Provides visualization of the traffic flow for a selected service. Select a service from the topology and click **Traffic Flow** from the right pane and select the clients that you want to monitor.
- View device information such as CPU and memory utilization, packets transmitted and received on the interfaces, logical resources and so on. Select a device in the topology to view device information.

14

CHAPTER

Assurance

IN THIS CHAPTER

- [Impact Analysis Overview | 176](#)
 - [About the Impact Analysis Page | 180](#)
 - [Key Anomalies in Impact Analysis | 181](#)
 - [Predictions Overview | 185](#)
 - [About the Predictions Page | 197](#)
-

Impact Analysis Overview

IN THIS SECTION

- [Impact Analysis Overview | 176](#)
- [Simulate Failure Scenarios | 178](#)
- [View Device Details | 178](#)
- [Benefits of Impact Analysis | 179](#)

Impact Analysis Overview

The Impact Analysis feature combines Apstra's intent-based analytics (IBA) with network service visibility. The Impact Analysis page on Juniper Data Center Assurance provides a visual representation of the network infrastructure with the following information:

- how nodes in the data center are connected to the clients
- what path the service traffic takes
- how the application traffic is handled
- how an anomaly in a node in the data center can impact application traffic.

Whenever an event occurs in the data center, Apstra generates several anomalies, which might include subsequent events triggered by an earlier anomaly. The main anomaly that triggered the related events is identified as the key anomaly. For example, when a port on a switch goes down, the event might generate various related anomalies like interface issues, cabling issues, route issues, and so on.

The Impact Analysis feature uses knowledge graphs to correlate events to issues in the data center network and then maps those issues to the potential impact on the applications that are utilizing the network. Instead of displaying several anomalies and causing alert fatigue, the Impact Analysis feature lets you view the key anomaly and the impacted devices in the topology diagram.

To view the Impact Analysis page, navigate to **Assurance > Impact Analysis**.

The **Anomalies** tab in the right pane of DC Assurance displays the current anomalies that occurred in the network along with their start time. Expand a key anomaly to view the detailed list of associated anomalies, symptoms, affected devices, interfaces, services, Virtual Networks, and VRFs. The tab also

displays a list of uncorrelated anomalies. Select an anomaly to highlight the affected node in the topology diagram.




NOTE: When you select a Link Status Mismatch anomaly, the affected node will not be highlighted on the topology diagram. It will only be listed in the anomaly details.

The **Changes** tab displays the changes in the network topology, such as addition or removal of a device. You can also use the filter on the page to examine how a specific service performed during the selected time period.

The Impact Analysis feature thus enables the administrators to quickly identify the root cause of an issue and resolve it so that the related anomalies also might get resolved. This helps network administrators reduce the mean time to resolution (MTTR), which in turn helps improve application user experience.

Figure 1 on page 177 shows the Impact Analysis page displaying how the nodes in the data center network are connected to the clients and how an anomaly in a data center device can impact the clients.

Figure 34: Impact Analysis Page




Video: [Impact Analysis](#)

Simulate Failure Scenarios

Using the Impact Analysis feature, you can also simulate failure scenarios to identify potential impact of an issue in the data center network. You just need to select a device in the network and click **Fail Scenario**. DC Assurance generates a visualization of how clients and services will be impacted if an issue occurs in the device. Thus, network administrators can use the Impact Analysis feature to analyze the impact of a network issue, and also to identify which component in the network is causing issues to the application traffic. Instead of providing generic solutions to troubleshoot issues, the Impact Analysis feature helps administrators to pinpoint specific issues in the data center network that affect application user experience.

Figure 2 on page 178 shows how an issue in the device apstra-fabric-001-leaf1 can impact the clients that are connected to it.

Figure 35: Simulate Failure Scenario




View Device Details

In addition to this information, you can also view details of specific nodes in the topology. Click a node to view details of the node in the right pane. The displayed information includes CPU and memory

utilization, along with details of transmitted and received packets on each interface of the device as shown in [Figure 3 on page 179](#).

Figure 36: Device Information



NOTE: You must import the DC-Probes pack from DC Hub on each blueprint in the Data Center Director instance to enable streaming of CPU and memory utilization data to Juniper Data Center Assurance. For information on pack installation, see [DC Hub Overview](#).

Benefits of Impact Analysis

The Impact Analysis page provides the following benefits:

- Enables administrators to map application issues to issues in the data center network, enabling faster resolution issues minimizing the impact on end user experience.
- Helps reduce the MTTR of data center issues, enhancing end user experience.

- Helps quickly identify the potential impact of network events on clients and services utilizing the network.

About the Impact Analysis Page

SUMMARY

The Impact Analysis page provides a graphical representation of the data center network along with visualization of how network issues can impact services and clients that are utilizing the data center network.

IN THIS SECTION

- [Tasks You Can Perform | 180](#)

The Impact Analysis page provides a graphical representation of the network infrastructure along with the path that the service traffic takes. The visualizations enable network administrators to quickly correlate anomalies to issues in the data center network and resolve them, thereby minimizing the impact on end users.

Tasks You Can Perform

- Filter and view impact analysis information for specific sites.
- View information about the services and clients that were impacted during different time intervals—Use the filter option to select the time interval. Move the slider at the bottom of the topology to view the changes during different time periods.
- Simulate failure scenarios—View potential impact on services and clients if a node in the data center network fails. Click **Fail Scenario** to view potential impact on clients in case of a node failure.
- View information for a selected site or for all the sites—Select a site to view information about a specific site or select **All Sites** to view information about all sites.
- View information about anomalies and changes in the network—The **Anomalies** Tab in the right pane displays the anomalies generated along with details of the services and clients that the anomaly impacts.
- View the potential impact of network issues on clients—Click **Ungroup** to view the potential impact of a network issues on traffic. The topology visualization provides how an issue with a node in the data center network can impact the clients that are connected to it.

- View memory and CPU utilization—View details of CPU and memory utilization information for the selected device for the current and previous week.
- View Tx and Rx packet count—View the count of packets transmitted and received through each interface on the selected device. You can also view the packet error count.

RELATED DOCUMENTATION

| [Impact Analysis Overview](#) | 176

Key Anomalies in Impact Analysis

SUMMARY

Understand the key anomalies listed in Impact Analysis.

The Impact Analysis feature displays anomalies that are detected within the Data Center network with the help of predefined IBA probes. On a Apstra Data Center Director instance, an IBA probe takes a set of data from the network, performs aggregations and calculations on it, and raises an anomaly if specific conditions are not met. This information is sent to the Juniper Apstra Edge device which further passes on the data to Juniper Data Center Assurance. The Impact Analysis feature determines the key anomaly and its associated anomalies from this data. It also pulls reports on the services and clients affected by the key anomaly and displays comprehensive information in the **Impact Analysis** page of the Juniper Data Center Assurance GUI. This feature enables the network administrator to quickly identify and resolve the root cause of an anomaly, thus reducing the impact on the network.



NOTE: You must enable IBA probes on the Apstra Data Center Director instance in order to detect anomalies in the network. For more information, see [Instantiate Predefined Probe](#) in the *Apstra Data Center Director User Guide*.

The tables below list key anomalies displayed in the **Anomalies** tab on the DC Assurance UI and the IBA probes that raise the anomalies detected in the network.

- [Table 1 on page 182](#) lists key anomalies that are detected on the devices in the network.

- [Table 2 on page 183](#) lists key anomalies that are detected on the interfaces of any given device in the network.
- [Table 3 on page 184](#) lists key anomalies related to the virtual devices in the network.
- [Table 4 on page 185](#) lists key anomalies that you can create and enable if required in your organization.

Table 44: System Key Anomalies

Key Anomaly	Description	IBA Probe
Fan Missing	Indicates that the number of fan trays for a device is less than expected.	Device Environmental Checks
Fan Status	Indicates that the fan state is not as expected.	Device Environmental Checks
Power Supply Missing	Indicates that the number of operational power supplies (PSUs) on a device is less than expected.	Device Environmental Checks
Power Supply Status	Indicates that the power supply status is not as expected.	Device Environmental Checks
Device CPU	Indicates that the CPU usage has exceeded the specified threshold for the specified duration.	Device System Health
Device Disk	Indicates that the disk usage has exceeded the specified threshold for the specified duration.	Device System Health
Device Memory	Indicates that the memory usage has exceeded the specified threshold for the specified duration.	Device System Health

Table 44: System Key Anomalies *(Continued)*

Key Anomaly	Description	IBA Probe
Drain Node	Indicates excess traffic on a node that is being drained.	Drain Traffic Anomaly

Table 45: Interface Key Anomalies

Key Anomaly	Description	IBA Probe
Cold Leaf Interface	Indicates that the traffic average for each spine-facing interface on a leaf is less than the acceptable range.	Hot/Cold Interface Counters (Fabric Interface)
Hot Leaf Interface	Indicates that the traffic average for each spine-facing interface on each leaf is higher than the acceptable range.	Hot/Cold Interface Counters (Fabric Interface)
Cold Specific Interface	Indicates that the average counter value of the specific interface is lower than the threshold value.	Hot/Cold Interface Counters (Specific Interface)
Hot Specific Interface	Indicates that the average counter value of the specific interface is higher than the threshold value.	Hot/Cold Interface Counters (Specific Interface)
Fabric Interface	Indicates that the interface has changed states (flapped) more than the configured threshold.	Interface Flapping (Fabric Interface)
Host Facing Interface	Indicates that the specific host-facing interface has changed states (flapped) more than the configured threshold number in recent history.	Interface Flapping (Specific Interface)

Table 45: Interface Key Anomalies (*Continued*)

Key Anomaly	Description	IBA Probe
Optical Transceiver	Indicates that the optical interface has crossed at least one threshold.	Optical Transceivers
Cabling Problem	Indicates an error with cabling.	Service Anomaly
L2 Interface to Server	Indicates a service anomaly in the L2 external server-facing interfaces.	Service Anomaly
L3 Interface to Server	Indicates a service anomaly in the L3 external server-facing interfaces.	Service Anomaly
Link Status Mismatch	Indicates a discrepancy in the link status.	Service Anomaly

Table 46: Key Anomalies on Virtual Devices

Key Anomaly	Description	Alarm
VM Resource Health Issue	Indicates that the CPU or memory utilization of the VM has exceeded the specified threshold for the specified duration.	Virtual Machine CPU Usage Virtual Machine Memory Usage
Network Connectivity Lost	Indicates that connectivity is lost between a VM and a host device. This anomaly is raised on the host device as an associated anomaly under Link Status Mismatch or as a key anomaly.	Network Connectivity Lost

Table 47: Additional Key Anomalies

Key Anomaly	Description	IBA Probe
Device Cold Anomalous	Indicates that the percentage of cold interfaces on a specific device is higher than the acceptable range.	Hot/Cold Interface Counters (Fabric Interface) Hot/Cold Interface Counters (Specific Interface)
Device Hot Anomalous	Indicates that the percentage of hot interfaces on a specific device is higher than the acceptable range.	Hot/Cold Interface Counters (Fabric Interface) Hot/Cold Interface Counters (Specific Interface)

Predictions Overview

SUMMARY

This page describes how the Predictive Analytics feature notifies administrators of potential issues in the network using machine learning models to predict future behaviors and outliers.

IN THIS SECTION

- [Predictions Overview \(Beta\) | 185](#)
- [System Health | 186](#)
- [Optical Interface Health | 187](#)
- [View Predicted Impacts | 190](#)
- [View Device Details and Service Topology | 194](#)
- [Benefits of Predictions | 197](#)


Predictions Overview (Beta)

Predictive Analytics uses machine learning (ML) techniques to forecast possible outliers in resource utilization patterns or equipment failures. This feature helps the administrator to proactively address issues and prevent possible outages.

The Apstra Edge device receives data about the network devices through the configured probes on Apstra. The Edge device sends this data to DC Assurance where the ML models perform the following tasks:

- Aggregate data points received from the network
- Train the ML models using these data points and learn normal behavior to create a baseline
- Forecast future data points by using trained ML models on historical data
- Identify deviations in the forecasted data when the future data points exceed the configured thresholds.

Figure 37: Predictive Analytics Workflow



The Predictive Analytics feature identifies outliers in the following metrics:

- System Health
- Optical Interface Health

System Health

The Predictive Analytics feature identifies outliers in the system CPU utilization and system memory utilization metrics. It uses the data collected by the Stream Receivers probe on Apstra to learn and forecast system health metrics and identify deviations. This feature allows you to proactively identify issues in the network infrastructure and resolve them before they impact network traffic.



NOTE: You can run the pre-flight check for the Edge device and verify whether the stream receivers are configured correctly. See [Pre-Flight Checks](#) for more information.

Predictive Analytics uses the following ML algorithms to predict System Health outliers:

Table 48: Machine Learning Algorithms for System Health

Function	Machine Learning Algorithm	Description
Forecasting	Light Gradient Boost Machine (LGBM)	An ensemble method based on decision trees that combines multiple weak models to produce a single strong prediction. Gradient boosting builds the model sequentially, with each model focusing on errors made by the previous one.



NOTE: Predictive Analytics uses ML algorithms to analyze the data it receives from the network to predict network behavior. Once this feature has been enabled, it requires a certain amount of data points to train its model on expected behavior and to recognize outliers. The accuracy and the confidence level of the predictions improve over time as more data is available for analysis.

Optical Interface Health



NOTE:

- Predictive Analytics for Optical Interface Health is a technology preview feature. For more information about technology previews, see [Juniper Apstra Tech Previews](#) in the *Data Center Director User Guide*.
- Predictive Analytics for Optical Interface Health requires your Data Center Director instance to be running version 6.1 or higher.
- Predictions for Optical Interface Health based on BER, FEC, and CRC metrics is possible only for devices running the following Junos OS Evolved releases: 23.4X100D31, 25.2X100D10, and 25.4R1.

Optical cables offer lower latency and better energy efficiency over long distances, which is crucial for maintaining performance across large-scale clusters. Optical interfaces are the ports that link the optical cables to devices for data transmission. An optical interface consists of multiple optical lanes that act as independent data transmission channels. Each lane carries a portion of the total data stream, and multiple lanes are aggregated to achieve higher overall bandwidth.

Identifying potential issues in the optical interfaces or lanes early enables network administrators to take corrective actions in advance, thereby preventing link outages and ensuring continuous network reliability.

The Predictive Analytics feature identifies outliers in the optical interfaces and lanes in the network. It uses the following data from Data Center Director to forecast optical health metrics and identify deviations from expected behavior:

- Digital Optical Monitoring (DOM) metrics data collected by the Optical Transceivers probe.
- Bit Error Rate (BER), Forward Error Correction (FEC), and Cyclic Redundancy Check (CRC) metrics data collected by a custom probe added in the DC-Probes analytics pack in DC Hub.



NOTE: To ensure that the Predictive Analytics feature can learn expected behavior and predict optical link health metrics,

- verify that the Optical Transceivers probe has been configured correctly on Data Center Director. For more information, see [Optical Transceivers](#) in the *Data Center Director User Guide*.
- import the DC-Probes pack from DC Hub on the site or blueprint. For information on pack installation, see [DC Hub Overview](#).

Predictive Analytics uses the following ML algorithm to predict Optical Interface Health outliers:

Table 49: Machine Learning Algorithm for Optical Interface Health

Function	Machine Learning Algorithm	Description
Forecasting	Multivariate Bidirectional Long Short-Term Memory (BiLSTM)	Bidirectional LSTM is a method that processes the input sequence in forward and backward direction, allowing the model to capture both past and future context of the input sequence. Multivariate BiLSTM processes and analyses data containing multiple related parameters simultaneously.

The Predictive Analytics feature uses DOM, BER, FEC, and CRC metrics data to predict optical metric trends and flag outliers indicative of link degradation.

[Table 50 on page 189](#) lists the DOM metrics and the signs that indicate degrading optical link.

[Table 51 on page 189](#) lists the BER, FEC, and CRC metrics and the signs that indicate degrading optical link.

Table 50: DOM Metrics Indicators of Optical Link Degradation

DOM Metrics	Description
Received optical power (Rx power)	Received optical power
Transmitted optical power (Tx power)	Output optical power
Laser bias current	Magnitude of the laser bias power setting current
Voltage	Voltage of the physical port
Temperature	Temperature of the physical port

Table 51: BER, FEC, and CRC Metrics Indicators of Optical Link Degradation

BER, FEC, and CRC Metrics	Description	Indicators of Link Degradation
FEC Corrected Codeword Rate	Number of errors corrected by FEC per second	Increasing values
FEC Uncorrected Codeword Rate	Number of errors that remained uncorrected by FEC per second	Increasing values
Pre-FEC BER	BER before FEC mechanism is applied	Increasing values
Rx CRC Errors Rate	Aggregated rate of increase in received CRC errors	Increasing values
Tx CRC Errors Rate	Aggregated rate of increase in transmitted CRC errors	Increasing values
FEC Histogram	Codeword percentage in each FEC histogram bins	Increasing values in the higher-error bins

Predictive Analytics for optical link health uses a statistical approach to detect outliers in the forecasted data. It marks outliers based on two factors:

- Outlier Detection by Trend

If the forecasted values of the RX power show a downward trend, the bias current and temperature show an upward trend for a specified number of consecutive days, an outlier will be raised.

- Outlier Detection based on Threshold

If the forecasted values of the optical health parameters cross the configured high alarm or moderate alarm thresholds for a specified duration, an outlier will be raised.




NOTE: Thresholds for the DOM metrics are received from the probe and configured based on the vendors of the optical transceivers and part number. The thresholds cannot be modified by a user.

View Predicted Impacts


To view the predictive outliers in the network, navigate to **Assurance > Predictions**. Use the site drop-down to select a specific site. You can also click the System Health or Optical Outliers cards to view outliers of the selected type.

Figure 38: Predictions Bubble Chart View



The **Predictions** tab displays a bubble chart of outliers mapped by expected time when the issue might occur and the confidence level. Each bubble represents an outlier predicted on a device within the network. The size of a bubble indicates the number of services impacted by that specific outlier. The color of the bubble indicates the severity level of the predicted outlier determined based on whether the forecasted value exceeds the system-defined thresholds. If the forecasted value exceeds the moderate threshold, an outlier with major severity level is raised. If the forecasted value exceeds the alarm threshold, an outlier with critical severity level is raised. The confidence level displays the reliability of the predictions in percentage. Select any bubble to view a summary of the predicted outlier. Click a bubble to view a summary card with information about the selected outlier.

Figure 39: Predicted Outlier Summary




NOTE: A single node in the network topology might have multiple predicted outliers displayed as overlapping bubbles. Different nodes in the network topology might also have multiple predicted outliers with similar confidence or severity levels displayed as overlapping bubbles. The summary card displays the total number of outliers overlapping the selected bubble. Click the bubble or the summary card to cycle through all the overlapping outliers.

You can also select the list icon on the top right of the Predictions section to view the predicted outliers in a list as shown in [Figure 40 on page 192](#). List View displays the predicted outliers in a table, along

with details of the affected device, the confidence and severity levels of the outlier, and the predicted time when the event might occur.


Figure 40: Predictions List View



NOTE: The Predictive Analytics feature forecasts system health outliers that might occur in a 24-hour period and optical interface health outliers that might occur in a 15-day period.


Use the drop-down options on the top right of the Predictions section to filter by the expected time when the issue might occur or by specific parameters. You can also use the **Filter** option to display predicted outliers based on whether services or clients are impacted, severity level of the outliers, or outliers predicted on specific devices in the network.

Figure 41: Predictions Filters



The right pane displays the number of impacted clients and services in the selected site. You can click the clients and services button to view and search from the full list of impacted clients.

Figure 42: Search Impacted Clients and Services



If no predictive outliers are displayed, it could be because of the following reasons:

- The network devices and optical links are behaving as expected.
- The Predictive Analytics feature is newly configured and not enough data is available yet to make reliable predictions.

- The streaming receivers and optical transceiver probes are not sending data from the Edge device to DC Assurance. In this case, verify whether the probes are configured correctly and that streaming is enabled on these probes on Apstra.

If the list of impacted clients and services for a predicted outlier is not displayed, it could be because:

- No traffic is flowing through the affected devices in the network.
- DC Assurance is not receiving traffic information from the Edge device. In this case, run the pre-flight check for the Edge device and verify if the flow servers are configured correctly. See [Pre-Flight Checks](#) for more information.

To view detailed information about an outlier,

1. Select a bubble from the bubble chart or an outlier from the Predictions list.
2. Click **View Details** from the summary card or select the **View Details** icon from the top right of the Predictions section.


The **Device Details** page opens with detailed metrics for the selected device.

View Device Details and Service Topology

On the **Device Details** page, graphs for the monitored parameters are displayed with predicted values for the selected device, along with a prediction band for each parameter. The prediction band indicates the range of values within which a future observation is likely to fall based on the confidence level of the forecasting method. A narrow prediction band indicates greater precision and confidence in the forecast. A wide prediction band indicates uncertainty and variability in the forecast.

An alert with all the outliers predicted for the selected resource is displayed at the top of the page along with the expected time of occurrence for each outlier.

Figure 43: Device Details page



Select the outlier category from the drop-down options on the top right to modify the displayed graphs as per your requirements.

- For System Health, the graph displays the historical, current, and predicted CPU and memory usage data for the selected device.
- For Optical Interface Health, select the optical interface from the drop-down and the page displays graphs with historical, current, and predicted values for all the optical interface health parameters for the selected device.

You can also choose the time interval for the data to be aggregated and displayed in the graphs from the drop-down options.

- For System Health, you can have the graph display data aggregated in 15-min or 1-hour intervals.
- For Optical Interface Health, you can have the graph display data aggregated in 1-hour, 12-hour, or 1-day intervals.


If an outlier is detected in the predicted data points, it is highlighted in purple on the graph. Hover over the outlier to view details such as the type of outlier, the severity, the predicted start and end time of the outlier, and so on.

Figure 44: Outlier Details



If the predicted outlier on the selected device is expected to have an impact on services and clients, the **Service Topology** option is enabled. Click **Service Topology** to view the network topology and traffic flow along with the impacted services and clients.

Figure 45: Outlier Topology



Benefits of Predictions

The Predictive Analytics feature provides the following benefits:

- Provides an early warning to administrators about possible failures.
- Helps understand the impact of predicted outliers.
- Enables administrators to proactively prevent potential outages.

About the Predictions Page

IN THIS SECTION

- [Tasks You Can Perform | 198](#)

The Predictions page provides a list of predicted outliers in the network that could cause outages and impact clients and services. This feature enables the administrators to quickly resolve potential issues in advance and prevent disruptions.

To access this page, click **Assurance** on the main menu and go to the **Predictions** tab.

Tasks You Can Perform

- Filter and view predictions information for specific sites.
- View predicted outliers for devices in the network along with the confidence level for each outlier displayed in the form of a bubble chart. The color of the bubble in the chart indicates the severity of the outlier. The size of the bubble indicates the number of services or clients impacted.
- View a summary of outliers predicted in the network.
- View a list of forecasted events in the network, their confidence level, their severity, and the predicted time of occurrence.
- View information about the services and clients impacted by the outliers.
- Select an outlier and view device metrics with graphs displaying historical, current, and predicted values for each parameter.
- View device metrics with data aggregated in the selected time intervals.
- View details about the predicted outliers in the metrics for the selected device.
- Select an outlier that impacts services or clients and view the topology connected to the device.

15

CHAPTER

Network Sustainability Analytics

IN THIS CHAPTER

- [Network Sustainability Analytics Overview | 201](#)
- [Forecast Power Cost and GHG Emission | 203](#)
- [Sustainability: Supported Platforms and Software Versions | 205](#)
- [Getting Started with Network Sustainability Analytics | 207](#)
- [Configure Sustainability Settings for the Organization | 213](#)
- [View Sustainability Insights for an Organization | 216](#)
- [Understanding the Organization-Level Sustainability Dashboard | 217](#)
- [Configure Sustainability Settings for a Site | 227](#)
- [View Sustainability Insights for a Site | 229](#)
- [Understanding the Site-Level Sustainability Dashboard | 231](#)
- [Review the Sustainability Report | 251](#)
- [View Sustainability Insights for a Device | 256](#)
- [Understanding the Device-Level Sustainability Dashboard | 258](#)
- [Configure Port Threshold for the Organization | 278](#)
- [Configure Port Threshold for a Site | 282](#)
- [Optimize Energy Consumption by Shutting Down Unused Ports | 285](#)
- [Optimize Low Traffic Ports Based on Historical Traffic Patterns | 289](#)
- [Optimize Low Traffic Ports Based on Forecasted Traffic Patterns | 293](#)
- [Optimize Energy Consumption by Shutting Down Unused Line Cards | 298](#)
- [Optimize Energy Consumption by Shutting Down Unused PFEs | 303](#)
- [Network Sustainability Analytics – Troubleshooting | 307](#)

Network Sustainability Analytics Overview

IN THIS SECTION

- [Benefits of Using the Network Sustainability Analytics Feature | 202](#)

As networks scale up exponentially to meet the rising bandwidth demands, energy consumption also goes up at a similar pace, substantially increasing the power required to run the networks uninterrupted. This leads to higher operational costs for organizations and contributes to a growing global carbon footprint.

Traditional network infrastructures often operate on an “always-on” model, where devices and components continue to consume power even during periods of low or no utilization. This results in inefficient energy usage and unnecessary costs.

In addition, limited visibility into real-time energy consumption and emission data makes it difficult for organizations to effectively monitor usage, identify inefficiencies, and align with sustainability goals.

To address these challenges, Data Center Assurance provides the Network Sustainability Analytics feature that enables organizations to monitor, analyze, and optimize energy consumption across their network infrastructure.

The Sustainability dashboards offer granular visibility into energy consumption, allowing organizations to track how devices and their components use energy across the network. It also provides granular insights into utilization patterns, helping identify idle or underutilized resources that can be safely powered down without impacting network performance.

The solution also offers predictive analytics capabilities. Using historical data, it forecasts future energy requirements, associated costs, and expected greenhouse gas emissions. These insights support informed decision-making for capacity planning, budgeting, and infrastructure upgrades.

In addition, the feature provides energy optimization recommendations that help transition from an “always-on” model to an “on-demand” network model, where network devices or components can be turned off during low-usage periods to conserve energy.

To further support decision-making, the potential energy savings feature estimates the reduction in energy consumption, emissions, and costs that can be achieved by applying optimization recommendations. This allows organizations to evaluate the impact of optimization actions before applying them.

The solution also helps enhance operational reliability by enabling faster detection and troubleshooting of power and thermal related anomalies. It helps identify devices having abnormal power and temperature fluctuations and helps you resolve issues proactively before they impact network performance.

In summary, the Network Sustainability Analytics feature in Data Center Assurance provides these capabilities:

- **Observability**—Enables you to monitor energy consumption and thermal profile, track GHG emission and power cost, and identify the top power consuming sites and devices through easy-to-use tables, charts, and graphs.
- **Insights on Sustainability Metrics**—Provides utilization trends and AI-generated forecasts based on historical data for SBTi compliance reports.
- **Energy Optimization Recommendations**—Provides intelligent recommendations to shut down unused or underutilized resources on the network to optimize energy consumption without impacting network performance.
- **Potential Energy Savings**—Provides estimated savings in energy, cost, and emissions to help evaluate the optimization potential of your network infrastructure.

Accessing Network Sustainability Analytics

The Data Center Assurance dashboard provides a high-level view of the sustainability analytics for the entire organization.

To view detailed sustainability analytics for the organization, site, or device, click **Sustainability** from the main menu and select the organization, site, or device using the scope selector.

To ensure accurate sustainability analytics, configure sustainability settings on the **Organization > Settings** or **Organization > Site Configuration** page by specifying the power cost, carbon intensity, and the port threshold values.

Benefits of Using the Network Sustainability Analytics Feature

- Gain detailed insights into sustainability metrics across organization, site, and device levels
- Enable data-driven sustainability planning and compliance reporting
- Evaluate potential savings in energy consumption, power cost, and emissions achievable across the network infrastructure
- Forecast the energy optimization potential of your network infrastructure to support budgeting and planning
- Reduce energy consumption by applying optimization recommendations

- Track greenhouse gas emissions and analyze trends to support sustainability goals
- Identify and troubleshoot power and thermal related anomalies across network devices

RELATED DOCUMENTATION

[Forecast Power Cost and GHG Emission | 203](#)

[View Sustainability Insights for an Organization | 216](#)

[View Sustainability Insights for a Site | 229](#)

[View Sustainability Insights for a Device | 256](#)

[Review the Sustainability Report | 251](#)

Forecast Power Cost and GHG Emission

IN THIS SECTION

- [Compliance with SBTi Targets | 203](#)
- [Power Consumption Forecast | 204](#)
- [Greenhouse Gas Emission Forecast | 204](#)

The Network Sustainability Analytics features in DC Assurance enable you to forecast energy consumption, power cost, and greenhouse gas emissions. To accurately measure and forecast power cost, you need to enter the power cost per unit provided by your power supplier in your preferred currency. For Carbon intensity data, you need to enter the emission value based on the source of power, regional emission standards, and your organization's emission goals.

Compliance with SBTi Targets


Science Based Targets initiative (SBTi) is a climate action organization that enables organizations to make their contribution in combating the climate crisis. SBTi targets are greenhouse gas emission reduction goals set by the organizations in line with Paris agreement on climate change.

SBTi target is a comprehensive score that includes energy consumption and carbon emission contributed by the entire infrastructure in your organization, which includes networking devices,

servers, power generation systems, cooling systems, and so on. DC Assurance focuses on providing sustainability analytics for network devices only, thereby enabling you to comply with the energy consumption targets and carbon emission standards for those devices. This data serves as an input for organization-level and site-level SBTi reports.

[Sustainability Report on page 204](#) shows the annual energy consumption, GHG emission, and power cost metrics. Hover over any of these values to view current value and the forecasted value for that metric.

Figure 46: Sustainability Report



Power Consumption Forecast

DC Assurance collects power usage data from the network devices and provides projections for energy consumption and power cost based on past consumption data. The forecasted data is useful for budgeting power cost. You can specify the power cost value in your preferred currency based on the source of power, power tariff rates in your geographical region, and your organization's contribution to green energy.

You can compare the estimated energy consumption and the actual energy consumed. If the actual power consumed exceeds the estimate, you can investigate the devices or components that are causing spikes in energy consumption, and then perform remedial actions.

Greenhouse Gas Emission Forecast

DC Assurance forecasts GHG emission, at the site-level and organization-level, based on the emission standards in your region and the power consumed by the network devices. You can specify a value based on the source of power and the emission standards in your geographical region. The total emission value correlates with the total energy consumed by the network devices in your organization.

Sustainability: Supported Platforms and Software Versions

SUMMARY

Use this topic to understand the minimum requirements for viewing Sustainability insights in Juniper Data Center Assurance.

To view Network Sustainability Analytics in Juniper Data Center Assurance, ensure that your deployment meets the following compatibility requirements.

[Table 1 on page 205](#) lists the support devices and the minimum required Junos OS versions.

[Table 2 on page 206](#) lists the minimum required Apstra Data Center Director and Apstra Edge versions.

Table 52: Supported Devices and Junos OS Versions

Supported Devices	Minimum Required Junos OS Version
EX4400-48T	22.4R3.25 23.4R2-S3.9
QFX5120-32C	22.2R3-S2.8 22.2R3-S4.10 23.2R1.13 23.4R2-S4.11
QFX5120-48Y-8C	22.2R3-S4.10 23.4R2-S4.11 23.4R2-S5.8
QFX5210-64C	23.4R2.13

Table 52: Supported Devices and Junos OS Versions *(Continued)*

Supported Devices	Minimum Required Junos OS Version
QFX5240-64QD	23.4R2-S2.3-EVO
QFX5130-48CM	23.4R2-S1.7-EVO
QFX5130-32CD	23.4R2-S3.11-EVO
QFX5230-64CD	23.4R2.14-EVO 23.4R2-S2.3-EVO
QFX5700	23.4R2-S2.3-EVO
PTX10001-36MR	23.4R2-S2.3-EVO
PTX10004	23.4R2-S2.3-EVO
PTX 10008	23.4R2-S2.3-EVO
PTX10016	23.4R2-S2.3-EVO

Table 53: Supported Apstra Data Center Director and Apstra Edge Versions

Apstra Data Center Director	Minimum Required Apstra Edge Version
4.2.1	0.45.5
5.0.0-64	
6.0.0-189	
6.1.0	

Getting Started with Network Sustainability Analytics

SUMMARY

Use this topic to get started with Network Sustainability Analytics in Juniper Data Center Assurance.

IN THIS SECTION

- [Access the Sustainability Feature | 207](#)
- [Complete the Initial Setup | 208](#)
- [Verify Prerequisites and Environment Readiness | 211](#)
- [Telemetry Accumulation and Refresh Cadence | 211](#)
- [Updates and Configuration Changes | 212](#)
- [Need Help? | 212](#)

To get started with Network Sustainability Analytics:

Access the Sustainability Feature

You must enable Network Sustainability Analytics feature for your organization.


The **Sustainability** menu is available in the sidebar.

When you click the **Sustainability** menu for the first time, the Sustainability landing page is displayed. However, you must configure the required sustainability parameters to view Sustainability data.

The landing page provides a preview of the key sustainability capabilities that will be available to you once the feature is enabled.

To begin using Network Sustainability Analytics, click the hyperlink on the landing page.

Figure 47: Sustainability Landing Page



Complete the Initial Setup

After you enable the Sustainability feature, complete the initial setup workflow.

- (Required) Keep the **Predictive Analytics** check box selected unless recommended otherwise by your HPE Networking representative.

When enabled, this feature offers AI-powered forecasts for key sustainability metrics, potential energy savings, and energy optimization recommendations.

- (Required) Configure the following mandatory inputs for calculations and reporting: **Average Power Cost (APC)** and **Carbon Intensity (CI)**.

For both parameters, suggested values are displayed based on global or location-specific estimates. You can click the provided hyperlink to automatically populate the suggested value in the corresponding field.

For more information, see ["Configure Sustainability Settings for the Organization" on page 213](#), ["Configure Sustainability Settings for a Site" on page 227](#).

- (Optional but Recommended) Configure the traffic thresholds for device ports by defining utilization ranges.

Use the **Port Traffic Thresholds** widget to define percentage ranges. Based on these ranges, the system classifies each port as **low**, **medium**, or **high** utilization.

The system then uses these utilization ranges to classify port behavior in [Per Port Peak / Off-Peak Traffic Monitoring on page 266](#) heatmap and to generate current and forecasted low traffic port optimization recommendations.

For more information, see ["Configure Port Threshold for the Organization" on page 278](#), ["Configure Port Threshold for a Site" on page 282](#).

Sustainability inputs can be configured at the organization and site levels. By default, Organization Settings is selected for configuration, see [Figure 2 on page 209](#).

To configure inputs for individual sites and obtain accurate, site-specific analytics, select **Site Customization**, see [Figure 3 on page 209](#).

Figure 48: Sustainability Settings at the Organization-Level

Sustainability

Please configure the APC (Average Power Cost) and CI (Carbon Intensity) at the Entire Org. Site.

☒ Organization Settings
☐ Site Customization

Organization Settings

☒ Predictive Analytics ⓘ
Carbon Intensity (gCO₂e/kWh) ⓘ
 ⓘ
[Suggested global average carbon intensity: 394 gCO₂e/kWh](#)
Average Power Cost (USD/kWh) ⓘ
 ⓘ
[Suggested global average power cost: 0.15 USD/kWh](#)
Currency ⓘ

Port Traffic Thresholds

0% 10% 80% 100%
Low Medium High

Figure 49: Sustainability Settings at the Site-Level

Sustainability

Please configure the APC (Average Power Cost) and CI (Carbon Intensity) at the Entire Org. Site.

☐ Organization Settings
☒ Site Customization

Site Customization
[Fill in Suggested CI and APC values for all sites](#)

Sites	Currency	APC (USD/kWh)	Suggested APC (USD/kWh)	CI (gCO ₂ eg/kWh)	Suggested CI (gCO ₂ eg/kWh)	Traffic Low %	Traffic High %
MainOrg-Site2	US Dollar	APC (USD/kWh) ⓘ	0.09	CI (gCO<sub>2</sub>eg/kWh) ⓘ	380	10	80
NEWSITE	US Dollar	APC (USD/kWh) ⓘ	0.09	CI (gCO<sub>2</sub>eg/kWh) ⓘ	380	10	80

NOTE:

- The currency unit at the site-level is inherited from the unit set at the organization-level and cannot be modified for individual sites.
- In the **Traffic Low %** and **Traffic High %** fields, enter utilization thresholds that define port utilization ranges in percentage.

Traffic Low % specifies the upper threshold for low port utilization.

Traffic High % specifies the lower threshold for high port utilization.

The system compares measured port traffic against these thresholds to classify ports as low or high utilization.

For example, if Traffic Low % is set to 30 and Traffic High % is set to 70, a port utilizing 20% of its capacity is classified as low utilization, while a port utilizing 85% of its capacity is classified as high utilization.

If not configured, the values are inherited from the organization-level settings.


You can update these settings at any time from the corresponding configuration pages. For more information, see ["Configure Sustainability Settings for the Organization" on page 213](#), ["Configure Port Threshold for the Organization" on page 278](#).

After the Sustainability feature is enabled and the required settings are configured, it may **take up to 24 hours** for Sustainability data to populate on the dashboards.

While data is being collected and processed, the Sustainability page displays a status banner indicating that information is still being collected and also shows a countdown until the page becomes active.

In the meantime, you can also view the Organization-level Sustainability dashboard layout in a greyed-out state. It provides a preview of the metrics, charts, and sections that will become available once data processing is complete. No actionable data is shown during this period.

Figure 50: Sustainability Dashboard Preview



Once the initial data processing completes, the Sustainability dashboard automatically becomes active and begins displaying sustainability analytics, insights, and recommendations.

Verify Prerequisites and Environment Readiness

Before Sustainability data can populate on the dashboards, ensure the following:

- The correct version of Apstra Edge is onboarded to Juniper Data Center Assurance.
Verify Edge status from the **DC Edges** page. See, "[Adopt Juniper Apstra Edge](#)" on page 124.
- Verify the devices and software versions meet the supported requirements to enable the Network Sustainability Analytics feature. See, "[Sustainability: Supported Platforms and Software Versions](#)" on page 205.

Telemetry Accumulation and Refresh Cadence

Sustainability dashboards may initially display partial or progressively populated data. Metric cards, trend charts, reports, and recommendations are updated as telemetry data accumulates and during subsequent refresh cycles.

Greenfield Deployments

- In a greenfield deployment, data collection begins after Apstra Edge is onboarded.
- Aggregated metrics and historical trends appear gradually as sufficient telemetry data is collected.
- A delay before meaningful data is visible is expected and normal and does not indicate a configuration issue.

Brownfield Deployments

- In a brownfield deployment, Apstra Edge is already onboarded or upgraded and telemetry collection is in progress.
- Sustainability dashboards populate more quickly by aggregating existing data.
- Aggregated metrics and historical trend charts may appear shortly after enablement.

**NOTE:**

- For both greenfield and brownfield deployments, it can take up to 24 hours for Sustainability data to appear on the dashboards.
- If dashboards remain empty after this period, verify Edge connectivity.

Updates and Configuration Changes

Once telemetry collection begins:

- Sustainability dashboards, reports, and recommendations are updated based on the standard data refresh cadence.
- Aggregated metrics, trend visualizations, forecasts, and recommendations reflect data collected in the most recent refresh cycle.
- Changes to sustainability settings, such as Average Power Cost, Carbon Intensity, or Port Traffic Thresholds, affect future calculations and recommendations only.
- Historical data is not recalculated when the sustainability settings are changed.

Need Help?

For additional assistance, contact your designated HPE Networking representative.

RELATED DOCUMENTATION

| [Network Sustainability Analytics – Troubleshooting](#) | 307

Configure Sustainability Settings for the Organization

Administrators can configure organization-wide carbon intensity and average power cost to ensure that the Organization-Level Sustainability Dashboard displays accurate GHG emissions and energy cost estimates.

The average power cost per kWh from the power supplier is used to calculate the aggregate power cost incurred by the organization.

The carbon intensity provided by the power supplier determines the GHG emissions of the organization. It represents the grams of CO₂e (carbon dioxide equivalent) emitted per kWh of energy consumed.

These inputs enable accurate environmental impact analysis and cost estimation for your organization.



NOTE:

- These inputs are mandatory when enabling the Sustainability feature for the organization. See, ["Getting Started with Network Sustainability Analytics" on page 207](#).
- Administrators can update these values at any time if the organization's carbon intensity or power cost changes.
- Administrators can configure these parameters both at the organization and the site level. If values are configured at both levels, the site-level configuration takes precedence over the organization-level configuration.

If these values are not configured, the dashboard may display default values or leave the fields blank. Configure these values to ensure that emission and cost calculations accurately reflect your organization's energy profile and support sustainability reporting, cost estimation, and long-term planning.



NOTE: Before you begin, ensure the following:

- You have the latest carbon intensity and power cost data from your power supplier or sustainability team.
- You have the required permissions to edit the **Organization Settings** page.

To configure sustainability settings for the organization:

1. Navigate to **Organization > Organization Settings**.

2. Locate the **Sustainability** tile.

Figure 51: Sustainability Tile

Sustainability

☒ Predictive Analytics ?

Carbon Intensity (gCO₂e/kWh) ?

394

[Suggested global average carbon intensity: 394 gCO₂e/kWh](#)

Average Power Cost (USD/kWh) ?

0.15

[Suggested global average power cost: 0.15 USD/kWh](#)

Currency ?

US Dollar

3. Keep the **Predictive Analytics** check box selected unless recommended otherwise by your HPE Networking representative.
When enabled, this feature offers AI-powered forecasts for key sustainability metrics, potential energy savings, and energy optimization recommendations.
4. To configure **Carbon Intensity**, enter the number of grams of CO₂e emitted per kilowatt-hour (gCO₂e/kWh) for the organization.
You must enter a non-negative numeric value, for example, **450**. The value can range from **0 to 1000**.
The system displays a suggested global average carbon intensity value for reference. To use this value, click the hyperlink to populate it in the field.
Carbon intensity varies by region, electricity provider, and the mix of renewable and non-renewable energy sources used to generate electricity.
Obtain this value from your power supplier, regional energy authority, or your organization's sustainability team.



NOTE: If supplier-specific values are not available, publicly available datasets, such as the EPA eGRID Data Explorer for US-based organizations, can be used for estimation purposes.

5. To configure the **Average Power Cost**, select your preferred currency from the **Currency** drop-down and enter your organization's average power cost per kilowatt-hour (kWh).



NOTE: The system supports a wide range of globally used currencies, allowing you to enter the cost in the currency that aligns with your regional, compliance, or financial reporting requirements.

If no currency is explicitly configured, the system defaults to **USD**.

Enter the amount you pay for a unit of electricity as a numeric value between **0.00** and **1.00**; for example, **0.12**.

The system displays a suggested global average power cost value for reference. To use this value, click the hyperlink to populate it in the field.

Use the electricity rate per unit negotiated with the power supplier.

If your organization uses variable or time-of-day pricing, use a blended average that is consistent with your organization's internal billing or reporting practices.

6. Click **Save** to apply the changes.



NOTE: Updates to sustainability settings are applied during the next data aggregation cycle and changes may take up to one hour to appear on the dashboards.

Once configured, the values are used to calculate the organization's current and potential GHG Emissions and Power Cost, which is displayed on the Organization-Level Sustainability Dashboard.

These values are also used to generate the organization-level energy report for SBTi planning and tracking and as inputs for aggregated calculations for any sites that do not have site-specific values configured.

To configure the sustainability metrics for specific sites, see ["Configure Sustainability Settings for a Site" on page 227](#). Site-level configurations take precedence over organization-level ones.



REMEMBER:

- Update average power cost whenever your electricity tariff changes.
- Update carbon intensity if your power supplier changes their energy mix or publishes new emission factors.

RELATED DOCUMENTATION

| [Network Sustainability Analytics – Troubleshooting](#) | 307

View Sustainability Insights for an Organization

The **Organization-Level Sustainability Dashboard** provides a high-level overview of your organization's overall energy performance. Use this dashboard to monitor key energy metrics, analyze historical trends, evaluate potential energy savings achievable through optimization actions across the organization, and identify high energy-consuming sites.

The dashboard is organized into two main tabs:

- **Overview**—Provides a high-level summary of your organization's energy profile, including energy consumption metrics for configurable time range, potential energy savings, trend charts, insights into top power-consuming sites, and carbon intensity and average power cost configured for the organization. See, "[Overview](#)" on page 217.
- **Potential Energy Savings**—Provides estimated energy, cost, and emissions savings achievable for the organization based on available device-level optimization recommendations. Use this tab to assess how much energy and cost your organization could save over various time ranges and to explore the recommendations, device types, and Junos OS versions that contribute most to estimated savings. See, "[Potential Energy Savings](#)" on page 225.

You can also access the organization-level Sustainability Report from this dashboard. See, "[Review the Sustainability Report](#)" on page 251.

Use the **Organization-Level Sustainability Dashboard** to:

- Monitor and understand organization-wide energy usage
- Identify high-impact opportunities for energy optimization
- Prioritize sites and devices that drive the highest energy consumption
- Estimate the sustainability benefits achievable through optimization actions

Understanding the Organization-Level Sustainability Dashboard

SUMMARY

Use the Organization-Level Sustainability Dashboard to monitor energy insights, potential energy savings, and top sites energy usage across the organization. The dashboard is organized into the two tabs: **Overview** and **Potential Energy Savings**.

IN THIS SECTION

- [Overview | 217](#)
- [Potential Energy Savings | 225](#)

Overview

The **Overview** tab provides a high-level summary of the organization's energy profile and highlights potential energy savings achievable through optimization actions across the organization.

You can view data for the selected time range. The available time ranges are **Today**, **Yesterday**, **Last 7 Days**, **Last 30 Days**, and **Custom**. By default, the widget displays data for **Last 30 Days**.




NOTE: The global time range selector applies to all widgets except the **Potential Energy Savings** widget, the **Power Usage** metric card, and the **Carbon Intensity** and **Average Power Cost** cards. The Potential Energy Savings widget displays projections for future time ranges, the Power Usage metric card shows the real-time power usage of the organization, and Carbon Intensity and Average Power Cost display configured values that do not change with the selected time range.

View Energy Metrics for the Organization

The organization-level metric cards display key aggregated energy metrics for the organization.

Use these summary cards to get a high-level overview of the organization's energy profile for the selected time range.

Figure 52: Summary Cards



For energy metric description, see [Table 54 on page 218](#).

Table 54: Energy Metrics Description

Metric	Description	Unit	Notes
Power Usage	Real-time power consumption of the organization, aggregated from site-level telemetry	Kilowatts (kW)	The timestamp indicates when the data was last updated.
Energy Consumption	Energy consumed by the organization during the selected time range	Kilowatt-hour (kWh)	The system aggregates power usage for the selected time interval to calculate energy consumption
GHG Emission	Total greenhouse gas emissions generated by the organization during the selected time range	gCO ₂ e	GHG emission is calculated based on the carbon intensity you configure at the organization level. If not configured, a default value is used to calculate the emission. See, "Configure Sustainability Settings for the Organization" on page 213.

Table 54: Energy Metrics Description (Continued)

Metric	Description	Unit	Notes
Cost	Estimated cost of energy consumed by the site during the selected time range	Region-specific currency	If the Total Cost appears blank, verify that the Average Power Cost is configured for the organization. See, "Configure Sustainability Settings for the Organization" on page 213.

**NOTE:**

- The timestamp in the top-right corner reflects the selected time range, and the values displayed for Energy Consumption, GHG Emissions, and Cost update based on this selection.
- The units displayed for Power Usage, Energy Consumption, and GHG Emissions are dynamically scaled based on the magnitude of the value.

View Historical Trend Graphs for the Organization

The **Historical Trend Chart** allows you to analyze energy trends for the organization over time.

You can view historical trends for the following metrics: Energy Consumption, GHG Emission, and Cost.

Click any of the three metric cards to view their trend charts for the selected time range in the right pane. When you select a time range, the trend graph displays data for the rolling interval ending at the most recent available data point.

Figure 53: Historical Trend Graph




Figure 2 on page 220 shows the organization's energy consumption trend during the last 30 days. You can view similar trend charts for the other two metrics as well.


The x-axis represents the selected time range and y-axis represents the aggregated organization-level values of the selected metric for the selected time range. Hover over any point on the graph to see the exact metric value for that timestamp.

View Potential Energy Savings for the Organization

The **Potential Energy Savings** widget provides a high-level summary of the projected energy savings, emissions reduction, and cost savings that can be achieved if all the available optimization recommendations are applied at the device-level.

The projections are displayed based on the selected time range. The available time ranges are **Today**, **Tomorrow**, **Next 7 Days**, and **Next 30 Days**. By default, the widget displays projections for **Next 30 Days**.

Figure 54: Potential Energy Savings Widget



For field descriptions, see Table 55 on page 221.

Table 55: Potential Energy Savings Widget

Energy Metric	Description	Data Availability
Projected Energy Savings	Estimated reduction in energy consumption that your organization can achieve by applying the available energy optimization recommendations.	The projection is displayed for the selected time range.
Reduction in GHG Emission	Estimated reduction in greenhouse gas emissions your organization can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and carbon intensity configured for the organization. NOTE: If the carbon intensity is not configured for the organization, a default value is used to calculate the reduction in GHG emissions. See, "Configure Sustainability Settings for the Organization" on page 213.	The projection is displayed for the selected time range.
Projected Cost Savings	Estimated cost savings that your organization can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and the average power cost configured for the organization. NOTE: If the Projected Cost Savings appear blank, verify that the Average Power Cost is configured for the organization. See, "Configure Sustainability Settings for the Organization" on page 213.	The projection is displayed for the selected time range.

Table 55: Potential Energy Savings Widget (*Continued*)

Energy Metric	Description	Data Availability
Potential Savings in %	Percentage of energy savings your organization could have achieved in the last 24 hours by applying all the available energy optimization recommendations. The value is calculated by comparing the potential energy savings over the last 24 hours with the total energy consumed during the same period.	The projection is displayed for the last 24 hours.

Use the ["Potential Energy Savings" on page 225](#) tab for deeper analysis of potential savings.

Navigate to the **Recommendation** tab on the Device-level Sustainability dashboard to review the available energy optimization suggestions.


View Carbon Intensity and Average Power Cost Configured for the Organization

The **Carbon Intensity** and **Average Power Cost** configured at the organization level are used to calculate GHG Emission and Cost displayed on the Organization-level Sustainability Dashboard.

The dashboard displays a banner prompting you to configure these values if they are not set.

You can configure these values on the Organization Settings page (**Organization > Organization Settings > Sustainability**). See, ["Configure Sustainability Settings for the Organization" on page 213](#).

Figure 55: Carbon Intensity and Average Power Cost Widgets



Carbon Intensity—Represents the amount of greenhouse gas emissions produced per unit of energy consumed. It is measured in grams of CO₂e per kWh (gCO₂e/kWh).

This value is used to calculate **GHG emissions** displayed on the dashboard.

GHG Emissions (CO₂e) = Energy Consumption (kWh) × Carbon Intensity (gCO₂e/kWh)

Configure an accurate carbon intensity value to ensure that the GHG emissions displayed on the dashboard reflect your organization's actual environmental impact.

Average Power Cost—Represents the cost per unit of electricity used at the organization-level.

This value is used to calculate the **Cost** displayed on the dashboard.

Configure this value to ensure accurate cost calculations across the dashboard.



NOTE:

- If carbon intensity is not configured, the default value of 443 will be used. You can override this value as needed.
- If average power cost is not configured, cost-related metrics, such as **Total Cost** and **Projected Cost Savings**, appear blank.
- The timestamps on the corresponding cards indicate when the values were last updated.

View Top Sites in the Organization By Energy Metrics

The **Top Sites by Energy Consumption** section provides a list of high energy consuming sites sorted in the decreasing order of energy consumption.

Figure 56: Top Sites by Energy Consumption

Top Sites by Energy Consumption ⓘ
From Feb 26, 2026 To Mar 28, 2026

Map View

Site Name	Energy Consumption	GHG Emission	Energy Cost	Traffic	Energy Savings
aos_...-DATA	7.35 MWh	1.75 tCO ₂ e	\$3,172.76	186.63 Gbps	48.42 MWh
aos_...-DATA	7.22 MWh	2.23 tCO ₂ e	\$4,144.99	30.30 Gbps	71.47 MWh
aos_...-AGMT	2.88 MWh	919.92 kgCO ₂ e	\$1,597.19	17.11 Mbps	404.96 MWh
aos_...-MGMT	2.70 MWh	779.55 kgCO ₂ e	\$1,444.45	70.69 Mbps	363.63 MWh
aos_...-POD	2.41 MWh	779.28 kgCO ₂ e	\$1,362.18	24.71 Tbps	239.81 MWh

Click the arrow next to any column header to sort the table based on that metric. The widget header updates accordingly to indicate the active sorting metric.

You can further drill down into a site to view detailed energy insights and identify devices contributing to higher energy consumption, cost, and GHG emissions.


Alternatively, you can use the scope selector at the top of the page to view sustainability insights for a specific site or a device.

View Top Sites on the Map

Toggle **Map View** on the Top Sites widget to view the geographical distribution of the sites on a map. The visualization on the map helps you quickly identify sites that have higher power consumption, cost, and GHG emissions.

Select a metric from the drop-down to update the map and compare sites. The size of the circle indicates how high or low the metric is at each site. Hover over a site to view detailed energy data for the past hour. You can also view the map in full screen mode and use pan or zoom to explore specific regions.

Figure 57: Map View




Potential Energy Savings

The **Potential Energy Savings** tab provides a high-level summary and a detailed breakdown of energy optimization potential across the organization. Use the tab to evaluate the potential energy savings achievable for the organization.

The tab includes the following widgets: **Potential Energy Savings** insights panel and **Average Hourly Savings** table.

The **Potential Energy Savings** insights panel provides aggregated savings across the organization over the selected time period. These insights are also available on the **Overview** tab. For information on the panel, see "[View Potential Energy Savings for the Organization](#)" on page 220.

Figure 58: Potential Energy Savings Tab



View Average Hourly Savings

The **Average Hourly Savings** table shows the average hourly savings achievable across the organization if all energy optimization recommendations are applied.

You can filter the table by recommendation type, device model, or Junos OS version installed on the device and get detailed breakdown of potential energy, emissions, and cost savings that contribute to the total potential savings shown on the insights panel.

Use the table to identify which recommendations, device types, or Junos OS version offers the largest potential energy savings.

Table 56: Average Hourly Savings Table

Field	Description
Energy Savings Category	Filter potential energy savings by one of the following categories: - Recommendation Type - Device Model - Device OS
Power Savings	The average hourly power savings for the selected category, assuming all applicable recommendations are applied.
CO2 Reduction	The average hourly reduction in greenhouse gas emissions for the selected category, assuming all applicable recommendations are applied.
Cost Savings	The average hourly cost savings resulting from power reductions for the selected category, assuming all applicable recommendations are applied.

Benefits of the Potential Energy Savings Feature

The Potential Energy Savings feature enables you to:

- Identify high-impact optimization opportunities and prioritize devices that require optimization
- Estimate the organization-level sustainability impact of applying energy optimization recommendations
- Assess the energy optimization potential of your network infrastructure

- Make data-driven decisions for improving energy efficiency and operational cost across the organization

Configure Sustainability Settings for a Site

Administrators can configure carbon intensity and average power cost for a site to ensure that the Site-Level Energy Dashboard displays accurate GHG emissions and power cost.

The average power cost per kWh from the power supplier is used to calculate the aggregate power cost for the site.

The carbon intensity provided by the power supplier determines the GHG emissions for the site. It represents the grams of CO₂e (carbon dioxide equivalent) emitted per kWh of energy produced.

You can configure these values both at organization and site levels. Site-level settings take precedence over organization-level settings. If values are not configured, the dashboard may show default values or leave fields blank.



NOTE:

- These inputs are mandatory when enabling the Sustainability feature for the organization. See, "[Getting Started with Network Sustainability Analytics](#)" on page 207.
- Administrators can update these values at any time if the site's carbon intensity or power cost changes.

Configure them to ensure that emission and cost calculations accurately align with the local power supplier's figures, reflect your site's energy profile and support sustainability reporting, cost estimation, and long-term planning.



NOTE: Before you begin, ensure the following:

- You have the latest carbon intensity and power cost data from your power supplier or sustainability team.
- You have the required permissions to edit the **Site Configuration** page.

To configure sustainability settings for a site:

1. Navigate to **Organization > Site Configuration**.
2. Click the site for which you want to configure these parameters.

3. Locate the **Sustainability** tile.

Figure 59: Sustainability Tile

Sustainability

Carbon Intensity (gCO₂e/kWh) ?

[Suggested value based on location: 700 gCO₂e/kWh](#)

Average Power Cost (USD/kWh) ?

[Suggested value based on location: 0.08 USD/kWh](#)

Currency ? (Set at org level: USD)



NOTE: For all the fields, enter numeric values only. You do not need to include the units in the input fields.

4. To configure **Carbon Intensity**, enter the number of grams of CO₂e emitted per kilowatt-hour (gCO₂e/kWh) for the site.

You must enter a non-negative numeric value, for example, **450**. The value can range from **0 to 1000**.

The system displays a suggested average carbon intensity value based on your location for reference. To use this value, click the hyperlink to populate it in the field.

Carbon intensity varies by region, power supplier, and the mix of renewable and non-renewable energy sources used to generate electricity for the site.

Obtain this value from your power supplier, regional energy authority, or your organization's sustainability team.



NOTE: If site or supplier-specific values are not available, publicly available datasets, such as the EPA eGRID Data Explorer for US-based sites, can be used for estimation purposes.

5. To configure the **Average Power Cost**, enter your site's average power cost per kilowatt-hour (kWh).



NOTE: The system supports a wide range of globally used currencies. The currency unit at the site-level is inherited from the unit set at the organization-level. See, "[Configure Sustainability Settings for the Organization](#)" on page 213.

Use the electricity rate per unit negotiated with your power supplier.

The system displays a suggested average power cost value based on your location for reference. To use this value, click the hyperlink to populate it in the field.

If your site uses variable or time-of-day pricing, use a blended average that is consistent with your organization's internal billing or reporting practices.

6. Click **Save** to apply the changes.



NOTE: Updates to sustainability settings are applied during the next data aggregation cycle and changes may take up to one hour to appear on the dashboards.

Once configured, the values are used to calculate the site's current and potential GHG emissions and power cost, which is displayed on the Site-Level Sustainability Dashboard. These values are also used to generate the site-level energy report for SBTi planning and tracking.



REMEMBER:

- Update average power cost whenever your electricity tariff changes.
- Update carbon intensity if your power supplier changes their energy mix or publishes new emission factors.

RELATED DOCUMENTATION

| [Network Sustainability Analytics – Troubleshooting](#) | 307

View Sustainability Insights for a Site

The Site-Level Sustainability Dashboard provides a comprehensive view of energy usage, efficiency trends, environmental impact, and potential energy savings achievable across a site.

The dashboard aggregates device-level telemetry from network devices and presents site-level insights through interactive visualizations, charts, and historical trend graphs.

In Juniper Data Center Assurance, a site represents a blueprint deployed in Apstra Data Center Director. When a blueprint is deployed, a site corresponding to the blueprint is created automatically in Juniper Data Center Assurance.

The site-level energy insights are organized into three tabs:

- **Overview** tab displays aggregated metrics and trend charts, potential energy savings achievable through optimization actions, and the top power consuming devices in the site. See, ["About the Overview Tab" on page 232](#).
- **Topology** tab provides a visual representation of device-level energy consumption patterns across the network fabric. See, ["About the Topology Tab" on page 243](#).
- **Potential Energy Savings** tab provides estimated energy, cost, and emissions savings achievable for the site based on available device-level optimization recommendations. See, ["About the Potential Energy Savings Tab" on page 249](#).

Use the Site-Level Sustainability Dashboard to:

Table 57: Top Tasks

If you want to	Use these resources
Analyze aggregated energy trends for a site, identify inefficiencies, anomalies, or abnormal power consumption patterns for the selected time range	"View a Site's Energy Metrics" on page 232
Analyze the potential energy, cost, emissions savings your site could achieve over various time ranges and identify the recommendations, device types, or Junos OS versions that contribute most to these estimated savings.	"About the Potential Energy Savings Tab" on page 249
Compare power and throughput of devices in a site to identify inefficient or underutilized devices	"Monitor Power and Throughput of Devices in a Site" on page 236
Analyze site efficiency using site-level power and throughput trends over time	"Monitor a Site's Power and Throughput Trends Over Time" on page 238
Identify top devices in the site across key energy dimensions, such as energy consumption, thermal deviations, and traffic patterns, and evaluate their impact on the site and plan optimization	"View the Top Devices in the Site by Selected Energy Metrics" on page 241

Table 57: Top Tasks *(Continued)*

If you want to	Use these resources
Visualize the site's network fabric and how devices compare in terms of key energy metrics	"About the Topology Tab" on page 243
Review a site's actual, historical, and forecasted energy data in the sustainability report to support your SBTi planning	"Review the Sustainability Report" on page 251
Adjust data views for more focused analysis. Filter data by time range or specific parameters to focus on the period or data type most relevant to your analysis	Use time range filters and metric selectors available across widgets

The Site-Level Sustainability Dashboard enables data center operators to make informed decisions that help improve energy efficiency, reduce operational costs, and support organizational sustainability goals.



ATTENTION: All values displayed on the dashboard are estimations based on the telemetry streamed every 5 to 15 minutes from the device. These values are intended for monitoring and analysis purposes and may not always reflect real-time measurements.

Understanding the Site-Level Sustainability Dashboard

SUMMARY

Use the Site-Level Sustainability Dashboard to monitor energy insights, identify inefficiencies, and gain visibility into site-level energy consumption. The site-level energy insights are organized into three tabs: **Overview**, **Topology** and **Potential Energy Savings**.

IN THIS SECTION

- [About the Overview Tab | 232](#)
- [About the Topology Tab | 243](#)
- [About the Potential Energy Savings Tab | 249](#)

About the Overview Tab

The Overview tab provides aggregated site-level energy metrics, trend visualizations, device energy insights, and potential energy savings to help you analyze the energy profile of your site.

For an overview of the Site-Level Sustainability dashboard, see ["View Sustainability Insights for a Site" on page 229](#).

You can view data for the selected time range. The available time ranges are **Today**, **Yesterday**, **Last 7 Days**, **Last 30 Days**, and **Custom**. By default, the widget displays data for **Last 30 Days**.



NOTE:

The global time range selector applies to all widgets except the **Potential Energy Savings** widget, the **Power Usage** metric card, and the **Carbon Intensity** and **Average Power Cost** cards. The Potential Energy Savings widget displays projections for future time ranges, the Power Usage metric card shows the real-time power usage of the organization, and Carbon Intensity and Average Power Cost display configured values that do not change with the selected time range.


A page-level scope selector also allows you to select a site. Based on the selected site, data is displayed on the energy dashboard.

View a Site's Energy Metrics

The site-level metric cards display key aggregated energy metrics for a site. These values are derived from device-level telemetry and aggregated at the site level.

Use the cards to get a high-level overview of the site's energy profile for the selected time range.

Figure 60: Metric Cards



For metrics description, see [Table 58 on page 233](#).

Table 58: Energy Metrics Description

Metric	Description	Unit	Notes
Power Usage	Real-time power consumption of the site, aggregated from device-level telemetry	Kilowatts (kW)	The timestamp indicates when the data was last updated.
Energy Consumption	Energy consumed by the site during the selected time range	Kilowatt-hour (kWh)	The system aggregates power usage for a rolling one hour interval to calculate energy consumption
GHG Emission	Total greenhouse gas emissions generated by the site during the selected time range	gCO ₂ e	GHG emission is calculated based on the carbon intensity you configure at the site or organization level. If not configured, a default value is used to calculate the emission. See, "Configure Sustainability Settings for a Site" on page 227 .
Cost	Estimated cost of energy consumed by the site during the selected time range	Region-specific currency	If the Total Cost appears blank, verify that the Average Power Cost is configured at the site or organization level. See, "Configure Sustainability Settings for a Site" on page 227 .

**NOTE:**

- The timestamp in the top-right corner reflects the selected time range for Energy Consumption, GHG Emissions, and Cost and the values update based on the selection.


- The units displayed for Power Usage, Energy Consumption, and GHG Emissions are dynamically scaled based on the magnitude of the value.

View Historical Trend Graphs for a Site

You can view historical trend graphs for Energy Consumption, GHG Emission, and Cost for the selected site.

Click any of the three metric cards to view their trend charts for the selected time range in the right pane. When you select a time range, the trend graph displays data for the rolling interval ending at the most recent available data point.

Figure 61: Historical Trend Graph



[Figure 2 on page 234](#) shows the site's emission trend for the last 30 days. You can view similar trend charts for the other two metrics as well.


The x-axis represents the selected time range and y-axis represents the aggregated site-level values of the selected metric for the selected time range. Hover over any point on the graph to see the exact metric value for that timestamp.

These trends help you identify recurring usage patterns, correlate energy usage with cost and emissions, validate the impact of optimization actions, and detect power consumption anomalies that may require further investigation or planning.

View Potential Energy Savings for a Site

The **Potential Energy Savings** widget provides a high-level summary of the estimated energy savings, emissions reduction, and cost savings that a site can achieve if all the available optimization recommendations are applied at the device-level.

Figure 62: Potential Energy Savings Widget



The projections are displayed based on the selected time range. The available time ranges are **Today**, **Tomorrow**, **Next 7 Days**, and **Next 30 Days**. By default, the widget displays projections for **Next 30 Days**.

For field descriptions, see [Table 59 on page 235](#).

Table 59: Potential Energy Savings Widget

Energy Metric	Description	Data Availability
Projected Energy Savings	Estimated reduction in energy consumption that the site can achieve by applying the available energy optimization recommendations.	The projection is displayed for the selected time range.
Reduction in GHG Emission	Estimated reduction in greenhouse gas emissions the site can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and carbon intensity configured for the organization. NOTE: If the carbon intensity is not configured for the site, a default value is used to calculate the reduction in GHG emissions. See, "Configure Sustainability Settings for a Site" on page 227.	The projection is displayed for the selected time range.

Table 59: Potential Energy Savings Widget (*Continued*)

Energy Metric	Description	Data Availability
Projected Cost Savings	Estimated cost savings that the site can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and the average power cost configured for the organization. NOTE: If the Projected Cost Savings appears blank, verify that the Average Power Cost is configured for the organization. See, "Configure Sustainability Settings for a Site" on page 227.	The projection is displayed for the selected time range.
Potential Savings in %	Percentage of energy savings the site could have achieved in the last 24 hours by applying all the available energy optimization recommendations The value is calculated by comparing the potential energy savings over the last 24 hours with the total energy consumed during the same period.	The projection is displayed for the last 24 hours.

Use the "[Potential Energy Savings](#)" on page 249 tab for deeper analysis of projected savings.


Navigate to the **Recommendation** tab on the Device-level Sustainability dashboard to review the available energy optimization suggestions.

Monitor Power and Throughput of Devices in a Site

The **Network: Power vs Throughput** scatter graph shows the relationship between throughput and power consumption for individual devices in the site during the selected time range.

Use this graph to visually compare how devices in the site differ in terms of the traffic they handle and the power they consume.

Figure 63: Network: Power vs Throughput scatter graph



Each circle represents a device within the site.

Table 60: Graph Description

Axis	Metric	Description	Unit
x-axis	Throughput	Traffic handled by the device during the selected time range. Throughput value in the graph is derived from the device's revenue ports (server-facing ports).	bps
y-axis	Power	Power consumed by the device during the selected time range.	Watts (W)



NOTE: Hover over a circle to view the device name and its actual throughput and power usage.

Understanding the position of each device on the graph helps you easily identify energy efficiency patterns.

- Bottom-right (high throughput, low power)**

These devices are delivering more traffic while consuming relatively less power. These are typically the most efficient devices in the site.
- Top-left (low throughput, high power)**

These devices are consuming more power than the traffic they handle. These devices may be inefficient or underutilized.

- **Upward diagonal trend**

These devices are consuming power proportional with throughput, which is typically the expected behavior for devices under traffic load.

- **Clusters**

Devices with similar power and throughput patterns.

Use the scatter graph to identify devices with high power usage but low throughput, detect underutilized or overloaded devices compared to their peers, and find opportunities for optimization.

Monitor a Site's Power and Throughput Trends Over Time

The **Power vs Throughput** line graph shows the aggregated power usage and throughput of all devices in the site during the selected time range.

Use the line graph to locate the site's inefficient time windows, when power consumption is high but throughput is low. You can view site-level trends across the selected time range.

Figure 64: Power vs Throughput Line Graph




Table 61: Graph Description

Axis	Metric	Description	Color	Unit
Left y-axis	Power	Shows the aggregated power usage of the site over the selected time range. NOTE: Click the Power legend to view only the power trend line on the graph. Double click the metric to clear the filter.	Brown	Watts (W)
Right y-axis	Throughput	Shows the aggregated traffic handled by the devices in the site over the selected time range. Traffic may vary significantly based on network activity. NOTE: Click the Throughput legend to view only the throughput trend line on the graph. Double click the metric to clear the filter.	Blue	bps
x-axis	Time Range	Represents the selected time range. By default, it is set to Last 30 Days.	—	Hours or Days



NOTE: Use the **Network: Power vs Throughput** scatter graph for device-level comparison and the **Power vs Throughput** line graph for site-level trend analysis.

Under normal operating conditions, power draw should increase when throughput increases and reduce during low-traffic periods. However, most network devices have a relatively flat power curve due to always-on components such as control planes, PSUs, transceivers, and cooling systems.

Use the chart to:

- Compare power to traffic proportionality in a site across time windows

- Identify idle power draw during off-peak hours and plan optimization efforts

For example, in a stock exchange network, throughput spikes during trading hours and drops afterward. However, power usage remains almost constant. You can use the historical trend graphs to identify off-peak hours and schedule port shutdown to avoid unnecessary power consumption.

- Compare site-level patterns across multiple sites


View Carbon Intensity Configured for a Site

Carbon intensity determines the site's GHG emissions. It represents the grams of CO₂e (carbon dioxide equivalent) emitted per kWh of energy consumed.

CO₂e is a standard unit used to express the global warming impact of all greenhouse gases in terms of their carbon dioxide equivalence.

The system uses carbon intensity as the input for calculating GHG emissions.

Figure 65: Carbon Intensity and GHG Emission Card



$$\text{GHG Emission (CO}_2\text{e)} = \text{Energy Consumption (kWh)} \times \text{Carbon Intensity (gCO}_2\text{e/kWh)}$$

Configure accurate carbon intensity value to ensure that the GHG emission shown on the dashboard reflects your site's actual environmental impact.

To configure carbon intensity for the site, see ["Configure Sustainability Settings for a Site"](#) on page 227.

If carbon intensity is not configured, a banner appears prompting you to configure this value. If left unconfigured, the system displays a default value sourced from recognized carbon data providers, which you can override as needed.




NOTE: The timestamp on the **Carbon Intensity** card indicates when the data was last updated. The **GHG Emissions** metric calculated using this value reflects emissions for the selected time range.

View the Average Power Cost Configured for a Site

The Average Power Cost card displays the cost per kWh of electricity used at the site.

This value is used to calculate the total power cost for the site. The average power cost is displayed in USD.

Figure 66: Average Power Cost and Total Cost Cards



To configure average power cost for the site, see, "[Configure Sustainability Settings for a Site](#)" on page 227.

If average power cost is not configured, a banner appears prompting you to configure this value. If left unconfigured, the Power Cost metric on the dashboard appears blank.

NOTE: The timestamp on the **Average Power Cost** card indicates when the data was last updated. The **Total Cost** metric calculated using this value reflects power cost for the selected time range.

View the Top Devices in the Site by Selected Energy Metrics

The Top Devices table lists all devices in the site and ranks them based on a selected energy metric, such as energy consumption, traffic, fan speed, or temperature deviation. The table aggregates device-level telemetry for the selected time range.

By default, the table lists devices in the descending order of energy consumption, allowing you to quickly identify the highest energy consuming devices in the site.

You can also sort the table based on other metrics available in the table. Click the arrow next to any column header to sort the table based on that metric.

For example, when you sort by **Temperature Deviation**, the table is reordered to display devices with the highest temperature deviation first. The widget header updates accordingly to **Top Devices by Temperature Deviation**, indicating the active sorting metric.

Figure 67: Top Devices Table

Top Device by Energy Consumption [®]						
From Feb 26, 2026 To Mar 28, 2026						
Device name	Energy Consumption	⚡ Max Traffic	Min Traffic	Average Fan Speed	Temperature Deviation	Energy Savings
voltus- 	579.14 kWh	886.5 Mbps	0 bps	3071.42 RPM	 -22.11 °C	16.50 kWh
qfx51- 	176.72 kWh	220.1 Mbps	0 bps	6374.29 RPM	 -25.07 °C	3.92 kWh
qfx- 	112.83 kWh	23.3 Gbps	1.6 kbps	4660.08 RPM	 -30.55 °C	3.78 kWh
qfx51- 	108.33 kWh	4.6 kbps	0 bps	4977.22 RPM	 -35.20 °C	782.30 Wh
voltus- 	83.44 kWh	616.27 bps	0 bps	10687.31 RPM	 -3.58 °C	112.20 Wh
qfx51- 	67.27 kWh	18.1 Gbps	0 bps	8433.12 RPM	 -14.51 °C	2.71 kWh
qfx51- 	63.91 kWh	8.6 Gbps	0 bps	11830.65 RPM	 -21.41 °C	3.23 kWh

For field description, see [Table 62 on page 242](#).

Table 62: Device Table

Field	Description
Device Name	Name of the device. Click the device name to navigate to the Device-Level Sustainability Dashboard for detailed insights.
Power Usage	The power utilized by the device during the selected time range, measured in Watts (W).
Energy Consumption	The total amount of energy consumed by the device during the selected time range, measured in kilowatt-hours (kWh).
Max Traffic	The maximum throughput recorded for the device during the selected time range, measured in bps.
Min Traffic	The minimum throughput recorded for the device during the selected time range, measured in bps.
Average Fan Speed	The average rotational speed of the device's fans during the selected time range, measured in Revolutions Per Minute (RPM).

Table 62: Device Table *(Continued)*

Field	Description
Temperature Deviation	The variation between the device’s measured temperature and the optimal temperature range. The icons help you quickly identify whether the device is running normal, warmer, or hotter than expected baseline. • Green—Indicates the device temperature is within expected range. No action required. • Yellow—The device is running slightly above its normal temperature but not yet critical. • Red—Temperature is significantly above normal and may impact device performance. Higher deviations may indicate thermal concerns such as airflow issues or fan failures.
Energy Savings	The energy savings the site could have achieved during the selected time range if all available energy optimization recommendations for the selected time range were applied.

Use the device table to:

- Identify devices that contribute the most to the site’s overall energy consumption
- Detect unusually high energy consuming devices
- Compare devices across multiple energy and thermal metrics
- Prioritize devices for optimization, reconfiguration, or replacement

About the Topology Tab

The Topology tab provides an interactive visualization of the site’s network fabric and uses color-coding to highlight device-level energy insights.


The topology is derived from the Data Center Director blueprint and reflects the layout of super spine, spine, and leaf devices in the site.



NOTE: You must adopt Apstra Edge to the organization to view topology and telemetry data of devices in the dashboard. See, "[Adopt Juniper Apstra Edge](#)" on page 124.

The topology view shows how devices in the network fabric are interconnected and how they compare in terms of power usage, energy consumption, and thermal behavior.

Figure 68: Topology View with Emissions Overlay



Key Elements in the View

- Each node represents a network device in the site.
- Device roles are color-coded (leaf, spine, super spine).

The legend in the bottom left corner indicates the role and corresponding color.

- Links represent physical or logical connectivity between devices.
- A color scale in the bottom right corner indicates the intensity range for the selected metric.

Use the metric selector on the top-right corner to overlay different metrics on the topology.

The available metrics are:

Table 63: Metric Overlay Options

Metric	Description	Color Gradient
Power	The device's power consumption.	Power in W 
Emissions	The device's estimated greenhouse gas emissions.	Emission in gCO₂e 
Temperature	The operating temperature of the device.	Temperature in °C  • Green—Operating within normal temperature range. • Yellow—Approaching defined temperature threshold. • Red—Exceeded defined temperature threshold.




NOTE: When a metric overlay is applied, devices are color-coded according to their value for that metric.

Darker or more intense shades indicate higher values, while lighter shades indicate lower values.

Use the color scale to quickly identify devices with higher energy or environmental impact and plan optimization or replacements.

You can use the [Juniper's Power Calculator](#) to compare power characteristics across device models.

Figure 69: Topology View



Topology View Task Pane

Table 64: Options on the Topology View

Options	Description
1	Metric Selector —Use the drop-down to choose a metric to overlay on the topology. The options are: - Emission - Power - Temperature
2	Reset —Reset the topology view to its default zoom position.
3	Layer Control —Switch between different data overlays.
4	Up a Level —Navigate up a topology hierarchy level.
5	Enter Fullscreen —Expand the topology view to fullscreen mode.
6	Refresh —Click to load the latest topology and telemetry data of the site.

Table 64: Options on the Topology View (Continued)

Options	Description
7	Color Scale —Shows the color range for selected metric. Darker or more intense shades indicate higher values, while lighter shades indicate lower values.
8	Device Legend —Maps device roles with their associated colors.

Interactive Actions

- Hover over a node to view device name, type, and the selected metric's value.
- Expand the leaf group in the topology view

By default, leaf nodes are grouped to simplify the topology view.

To view individual leaf nodes:

1. Click the leaf group node to open the search and selection panel.

Figure 70: Search and Selection Panel




2. Use the drop-down to select the IP address of the leaf nodes you want to display in the topology view.

Click **Select All** to display all leaf nodes.

3. Click the checkmark to apply your selection.

Figure 71: Topology View Navigation



Use the navigation arrows in the topology view to pan across the fabric and explore different sections of the layout.

- Click a node to open the knowledge card.
The card provides more details on the device, such as traffic details and detected anomalies.
- Double click a device to highlight only its connections.
Click **Reset** to reset the view.

Benefits of the Topology View

Use the topology view to:

- Visualize the layout of the site's network fabric
- Identify high power consuming or inefficient devices in the context of their fabric role
- Identify hotspots such as high power consuming or overheated devices.
- Detect tier-level patterns in energy usage. For example, a single spine device consuming significantly more power than its peers
- Correlate energy behavior with device roles (leaf, spine, super spine) to identify anomalies such as a single device consuming far more power than other devices in the site


About the Potential Energy Savings Tab

The **Potential Energy Savings** tab provides a high-level summary and a detailed breakdown of energy optimization potential for a site. Use this tab to evaluate the optimization and savings potential before applying the recommendations.

The tab includes the following widgets: **Potential Energy Savings** insights panel and **Average Hourly Savings** table.

The **Potential Energy Savings** banner provides aggregated savings for the site over the selected time period. These insights are also available on the **Overview** tab. For information on the panel, see ["View Potential Energy Savings for a Site" on page 234](#).

Figure 72: Potential Energy Savings Widget



View Average Hourly Savings

The **Average Hourly Savings** table shows the average hourly savings achievable for a site if all energy optimization recommendations are applied.

You can filter the table by recommendation type, device model, or Junos OS version installed on the device and get detailed breakdown of potential energy, emissions, and cost savings that contribute to the total projected savings.

Use the table to identify which recommendations or device types offer the largest potential energy savings in the site.

Table 65: Average Hourly Savings Table

Field	Description
Energy Savings Category	Filter potential energy savings by one of the following categories: • Recommendation Type • Device Model • Device OS
Power Savings	The average hourly power savings for the selected category, assuming all applicable recommendations are applied.
CO2 Reduction	The average hourly reduction in greenhouse gas emissions for the selected category, assuming all applicable recommendations are applied.
Cost Savings	The average hourly cost savings resulting from power reductions for the selected category, assuming all applicable recommendations are applied.

Navigate to the **Recommendation** tab on the Device-level Sustainability dashboard to review the available energy optimization suggestions.

Benefits of the Potential Energy Savings Feature

The Potential Energy Savings feature enables you to:

- Identify high-impact optimization opportunities within the site and prioritize devices that require optimization
- Estimate the site-level sustainability impact by evaluating the potential energy savings
- Assess the overall energy optimization potential of the site
- Compare energy optimization potential across different sites in your organization
- Make data-driven decisions for improving energy efficiency and reducing operational cost of the site

Review the Sustainability Report

SUMMARY

Use this topic to understand how to evaluate sustainability reports at both the organizational level and at an individual site level.

IN THIS SECTION

- [Benefits of Using the Sustainability Report | 255](#)

The Sustainability Report provides the visibility required to plan, track, and meet your organization's Science Based Targets initiative (SBTi) goals. The report summarizes actual and historical energy performance and uses AI-based forecasting to project future usage, energy cost, and emissions for the selected time range, supporting flexible, time-based analysis. By default, the report displays data for the current calendar year.

The report is available at both the organization and the individual site levels. The organization-level report is aggregated from all the sites in the organization.

By combining actual usage, cumulative yearly metrics, historical trends, and AI-based forecasts, the Sustainability Report helps support long-term sustainability planning and reporting at both the organization and site levels.

The Sustainability Report distinguishes between two types of greenhouse gas emissions:

Operational Emissions—Emissions generated from energy consumption during device operation. These emissions vary over time based on network activity and workload. The **GHG Emissions** metric card and corresponding trend chart represent operational emissions.

Non-Usage Emissions—Emissions associated with the lifecycle of network devices, including manufacturing, transportation, and end-of-life disposal. These emissions are independent of day-to-day device operation and are reported separately in the **Emissions from Non-Usage Product Lifecycle Phases** section of the report.

Together, operational and non-usage emissions provide a more complete view of the carbon footprint. Reviewing both metrics helps support sustainability planning, infrastructure lifecycle management, and hardware procurement decisions.

When reviewed together, the organization-level and site-level sustainability reports provide a clearer view of your network's overall energy profile. The organization-level report highlights high-level trends, while the site-level reports provide granular visibility into the underlying factors. This combined view helps identify where corrective actions are required and whether they should be addressed at an individual site-level or across multiple locations in the organization.

To review Sustainability Report for the organization, click the **Sustainability Report** icon in the top-right corner of the **Organization-Level Sustainability Dashboard** to open the organization's sustainability report.

To review Sustainability Report for a site, click the **Sustainability Report** icon in the top-right corner of the site's **Site-Level Sustainability Dashboard** to open the site's sustainability report.



NOTE: The reports are refreshed daily.

Click the **Download** icon in the top-right corner to export the report in CSV format. This allows you to integrate the data with internal compliance systems and to customize the data to suit your organization's specific reporting needs.

The Sustainability Report is organized into the following sections:

Metric Cards

The metric cards display the cumulative expected values for energy consumption, GHG emissions, and energy cost for the selected time range.



NOTE: The GHG Emissions card represents the operational emissions of across the organization or sites. To view the non-operational emissions check the **Emissions from Non-Usage Product Lifecycle Phases** section in the report.


Hover over a metric card to view the breakdown of actual and forecasted values that make up the cumulative total.



NOTE: Actual values are shown only for past periods. If the selected time range includes future months, only forecasted values are displayed for those months.

[Figure 1 on page 253](#) shows the cumulative GHG emissions and the breakdown of actual and forecasted values for the current calendar year.

Figure 73: Sustainability Report




Click a metric card to expand the corresponding accordion to view the monthly trend charts for the selected time range.

Trend Charts

The trend chart provides a month-by-month view of the historical and forecasted values for the selected energy metric (Energy Consumption, GHG Emission, and Energy Cost) for the selected time range. [Figure 2 on page 254](#) shows the yearly historical and forecasted energy consumption trends.

Figure 74: Trend Chart for Energy Cost



The months are listed on the left and the selected metric's value is listed on the right.

Each row contains a horizontal bar on which the historical data appears in grey and forecasted data in blue. Hover over the bar to view the individual historical and forecasted values for that month. The value displayed on the right represents the cumulative total of the historical and forecasted values.

Time periods without data appear blank.



CAUTION: Forecasted values are AI-generated estimates based on available historical data. While efforts have been made to ensure reliability, accuracy may vary depending on workload shifts, configuration changes, and site-level factors.

Non-Usage Emissions


The Sustainability Report includes non-usage emissions, which represent lifecycle greenhouse gas emissions associated with network devices. These include emissions from manufacturing, transportation, and end-of-life disposal, and are independent of operational energy consumption.

Expand the **Non-Usage Emissions** section to view detailed contributions from each phase. The report displays the total non-usage emissions in kgCO₂e and breakdown of emissions by lifecycle phase.

Time Range Selection

The Sustainability Report includes a **Time Range Selector** that allows you to view data across different periods.

Figure 75: Time Range Selector



You can select **This Year**, **Last Year**, **Next Year**, or define a **Custom Range**. By default, the report displays data for **This Year**.



NOTE:

- **This Year** displays data for the full calendar year, from January to December.
- **Next Year** appears in the drop-down after the current date passes the midpoint of the current year.
- **Last Year** appears in the drop-down only if historical data is available through January of the current calendar year.
- When selecting a custom date range, choose a minimum of 3 months to ensure your sustainability report contains meaningful data.

The selected time range applies to all metric cards and trend charts within the report. This enables analyzing rolling twelve-month performance, comparing year-over-year results, tracking improvements across reporting cycles, and gaining seamless visibility into historical and projected performance.

Benefits of Using the Sustainability Report

Use the report to:

- Visualize and review the actual and forecasted GHG emissions, energy consumption, energy cost for a selected time range. Flexible, time-based reports allow alignment with internal sustainability reporting cycles
- Export the report in CSV format to integrate with your organization's reporting needs
- Identify annual trends and periods of high energy use

- Anticipate potential increases in energy consumption, cost, or emissions using forecasted values
- Assess alignment with SBTi goals using projected performance
- Prioritize energy optimization and sustainability initiatives using data-driven insights
- Understand total lifecycle environmental impact, including both operational and non-usage emissions
- Share insights with internal sustainability teams, auditors, or stakeholders

View Sustainability Insights for a Device

In data center networks, traffic patterns can fluctuate significantly depending on workload. However, most network devices typically draw a near-constant level of power, even during low traffic windows. Always-on components, such as the forwarding plane ASICs, PSUs, transceivers, PHY devices, and cooling systems, contribute to this idle power draw and lead to substantial energy wastage during low utilization periods. To optimize energy usage, you need visibility at the individual device level.

The Device-Level Sustainability Dashboard provides you comprehensive sustainability analytics at the individual device level. You gain granular visibility into the device's energy consumption, how efficiently it operates under different traffic conditions, and how you can optimize power usage without impacting device performance.

Use the dashboard to monitor power usage, analyze power-to-traffic proportionality, review current and forecasted port activity, monitor thermal performance, and identify opportunities to reduce idle power draw without impacting device performance. You also receive intelligent, data-driven recommendations to guide energy optimization efforts across your data center and get to evaluate the potential energy savings achievable by applying the available recommendations.

The device-level sustainability analytics are organized into four tabs:

Table 66: Top Tasks by Tab

Tab	Top Tasks
"Overview" on page 258	• Monitor a device's power usage, energy consumption, GHG emission, power cost, and track historical trends • Evaluate potential energy savings possible for the device • Analyze a device's power vs throughput behavior • Review current and forecasted port peak and off-peak traffic patterns • Correlate port activity with power usage
"Power Supply" on page 268	Diagnose power supply efficiency, redundancy, and input-feed issues. • Verify device power capacity and utilization • Check power module status • Review detailed power module telemetry
"Thermal Profile" on page 271	Proactively manage thermal performance. • Monitor a device's temperature and fan speed trends • Analyze the device power trend • Review a device's top sensors by temperature deviation • Inspect fan status per tray
"Recommendation" on page 276	• Review energy optimization recommendations that can help optimize the device's energy consumption and improve sustainability • Evaluate the potential energy savings achievable by applying these recommendations

Each tab provides key sustainability analytics, time-series trends, and actionable insights based on device telemetry.



ATTENTION: All values displayed on the dashboard are estimations based on the telemetry streamed every 5 to 15 minutes from the device. These values are intended for monitoring and analysis purposes and may not always reflect real-time measurements.

Understanding the Device-Level Sustainability Dashboard

SUMMARY

Use the Device-Level Sustainability Dashboard to gain visibility into how much power a network device in your Data Center Director-managed data center consumes, how efficiently it operates under different traffic conditions, and where energy optimization is possible without impacting device performance. The sustainability analytics are organized into four tabs: **Overview**, **Power Supply**, **Thermal Profile**, and **Recommendation**.

IN THIS SECTION

- [About the Overview Tab | 258](#)
- [About the Power Supply Tab | 268](#)
- [About the Thermal Profile Tab | 271](#)
- [About the Recommendation Tab | 276](#)

About the Overview Tab

The **Overview** tab provides a high-level summary of a device's energy usage, efficiency, port activity, and potential energy savings.

Use this tab to quickly assess device's key energy metrics, analyze power-throughput behavior, identify ports contributing to idle power consumption, and evaluate the potential energy savings achievable by applying the recommended actions.

- **View Key Energy Metrics for a Device**

The summary cards display the key energy metrics for the device.

Use these cards to monitor how much power the device consumes, how that translates into energy cost and greenhouse gas (GHG) emissions, and how usage trends vary over time.

Figure 76: Device-level Summary Cards




Table 67: Energy Metrics Description

Metric	Description	Unit	Notes
Power Usage	Real-time power draw of the device	Kilowatts (kW)	The timestamp indicates when the data was last updated.
Energy Consumption	Energy consumed by the device during the selected time range	Kilowatt-hour (kWh)	The system aggregates power usage for a rolling one hour interval to calculate energy consumption
GHG Emission	Estimated greenhouse gas emission generated by the device during the selected time range	gCO ₂ e (grams of CO ₂ equivalent)	Calculated based on the carbon intensity you configure at the organization or site level. If not configured, a default value is used to calculate the emission. See, " Configure Sustainability Settings for the Organization " on page 213 or " Configure Sustainability Settings for a Site " on page 227.

Table 67: Energy Metrics Description (Continued)

Metric	Description	Unit	Notes
Cost	Estimated cost of energy consumed by the device during the selected time range	Region-specific currency	If cost per unit is not configured at the organization or site level, this field appears blank. See, "Configure Sustainability Settings for the Organization" on page 213 or "Configure Sustainability Settings for a Site" on page 227 .



NOTE:


- The values displayed for Energy Consumption, GHG Emissions, and Cost represent aggregated values for the selected time range. The timestamp indicates the selected time range. By default, the data is displayed for the last 30 days.
- The units displayed for Power Usage, Energy Consumption, and GHG Emissions are dynamically scaled based on the magnitude of the value.

Trend graphs for Energy Consumption, GHG Emission, and Cost are available for the device.

Click any of these three summary cards to view their historical trend for the selected time range in the right pane. The x-axis represents the selected time range and y-axis represents the selected metric's values. Hover over any point on the graph to see the exact metric value for that timestamp.

[Figure 2 on page 261](#) shows the power cost trend for the last 30 days. Hover over the trend line to see the cost for a specific timestamp in a popover.

Figure 77: Cost Trend Graph



- **View Potential Energy Savings For the Device**

The **Potential Energy Savings** widget provides a high-level summary of the projected energy savings, reduction in emission, and cost savings that a device can achieve if all the available optimization recommendations are applied.

Figure 78: Potential Energy Savings Widget

Potential Energy Savings				Next 30 days ▾
Projected Energy Savings	Reduction in GHG Emission	Projected Cost Savings	Potential Savings in %	
110.64 kWh	49.12 kgCO ₂ e	\$88.51	3.22%	

The projections are displayed based on the selected time range. The available time ranges are **Today**, **Tomorrow**, **Next 7 Days**, and **Next 30 Days**. By default, the widget displays projections for **Next 30 Days**.

For field descriptions, see [Table 68 on page 261](#).

Table 68: Potential Energy Savings Widget

Energy Metric	Description	Data Availability
Projected Energy Savings	Estimated reduction in energy consumption that the device can achieve by applying the available energy optimization recommendations.	The projection is displayed for the selected time range.

Table 68: Potential Energy Savings Widget (*Continued*)

Energy Metric	Description	Data Availability
Reduction in GHG Emission	Estimated reduction in greenhouse gas emissions the device can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and carbon intensity configured at the organization or site level. NOTE: If the carbon intensity is not configured at the organization or site level, a default value is used to calculate the reduction in GHG emissions. See, "Configure Sustainability Settings for the Organization" on page 213 or "Configure Sustainability Settings for a Site" on page 227.	The projection is displayed for the selected time range.
Projected Cost Savings	Estimated cost savings that the device can achieve by applying the available energy optimization recommendations. The value is calculated based on the projected energy savings and the average power cost configured at the organization or site level. NOTE: If the Projected Cost Savings appears blank, verify that the Average Power Cost is configured at the organization or site level. See, "Configure Sustainability Settings for the Organization" on page 213 or "Configure Sustainability Settings for a Site" on page 227.	The projection is displayed for the selected time range.

Table 68: Potential Energy Savings Widget *(Continued)*

Energy Metric	Description	Data Availability
Potential Savings in %	Percentage of energy savings the device could have achieved in the last 24 hours by applying all the available energy optimization recommendations. The value is calculated by comparing the potential energy savings over the last 24 hours with the total energy consumed during the same period.	The projection is displayed for the last 24 hours.

Navigate to the **Recommendation** tab to review the available energy optimization suggestions and the potential energy savings for each recommendation. See, ["About the Recommendation Tab"](#) on [page 276](#).

- **Monitor a Device's Power vs Throughput Patterns**

The **Power vs Throughput** line graph visualizes the historical correlation between a device's power usage and its throughput over a selected time range.

Use this graph to analyze how efficiently the device converts traffic load into power usage, identify idle power consumption, analyze historical patterns, such as, peak vs off-peak behavior, and identify opportunities for energy optimization.

Figure 79: Power vs Throughput Graph



Table 69: Graph Description

Axis	Metric	Description	Color	Unit
Left y-axis	Power	Shows the aggregated power usage of the device over the selected time range. NOTE: Click the Power legend to view only the power trend line on the graph. Double click the metric to clear the filter.	Brown	Watts (W)
Right y-axis	Throughput	Shows the aggregated traffic handled by the device over the selected time range. Traffic may vary significantly based on network activity. NOTE: Click the Throughput legend to view only the throughput trend line on the graph. Double click the metric to clear the filter.	Blue	bps
x-axis	Time Range (Today, Yesterday, Last 7 Days, Last 30 Days, Custom)	Displays the selected time range. By default, it is set to Last 30 Days.	—	Hours or Days

**NOTE:**


- The metric legends available below the graph acts as filters. Click a metric to view only its trend line on the graph. Double click the metric to clear the filter.

By default, the trend lines for both metrics are displayed on the graph.

- Hover over any point on the lines to view the exact power and throughput values for that timestamp.

Under normal operating conditions, power draw increases when throughput increases and reduces during low-traffic periods. However, most network devices have a relatively flat power curve due to always-on components such as forwarding planes, PSUs, transceivers, and cooling systems, as shown in "Figure 5" on page 265.

Figure 80: Sample Power vs Throughput Graph



For example, in a stock exchange network, throughput spikes during trading hours and drops afterward. However, power usage remains almost constant. You can use the historical trends to schedule port shutdown during off-peak hours to reduce unnecessary power consumption.

Use the graph to:

- Compare the device's power to traffic proportionality across time ranges
- Identify idle power draw during low-traffic periods and plan optimization efforts
- Compare device-level energy patterns across multiple devices
- **Monitor a Device's Per Port Peak and Off-Peak Traffic Patterns**

The **Per Port Peak / Off-Peak Traffic Monitoring** heatmap shows the hourly traffic pattern for each port on the device for a rolling 24 hour window.



NOTE: The traffic is categorized based on the port utilization thresholds configured at the organization or site level. For more information, see "[Configure Port Threshold for the Organization](#)" on page 278, "[Configure Port Threshold for a Site](#)" on page 282.

Use this visualization to easily identify unused or lightly used ports that contribute to idle power consumption.

For devices that support multiple line cards, you can view hourly traffic pattern per line card.




NOTE:

Currently, we support line card level traffic pattern visualization for the following platforms: PTX10016, PTX 10008, PTX10004, PTX10001-36MR.

Select the **Forecast** checkbox to display hourly predicted traffic patterns for the next 24 hours.

Figure 6 on page 266 shows the heatmap with the **Forecast** option enabled. The heatmap is split into two sections: current traffic pattern on the left and the forecasted traffic pattern on the right.

Figure 81: Per Port-Peak/ Off-Peak Traffic Heatmap



Forecasts are generated based on several days of historical port traffic data. Use the forecast to identify ports that show stable patterns of low utilization windows and schedule temporary port shutdowns.

Heatmap Components

- Each row corresponds to a port on the device. The count varies by device model.
- On modular devices with multiple line cards, ports are grouped by line card.
- Each cell represents an hourly time slot for a rolling 24 hour window.
- Hover over a tile to view the port number and traffic for that hour.
- The timestamp above the heatmap shows the exact 24 hour time window for which the data is displayed.

When forecasting is enabled, timestamps are shown for both the current traffic and forecasted traffic.

- Traffic intensity on each port is represented using gradients of green. See Table 70 on page 267



NOTE: Traffic on the heatmap is categorized based on the port traffic thresholds you configure at the organization or site level.

For more information, see ["Configure Port Threshold for the Organization" on page 278](#), ["Configure Port Threshold for a Site" on page 282](#).

Table 70: Traffic Intensity Based on Configured Port Traffic Threshold

Color	Traffic Level	Description
Dark Green	High	Port is handling heavy traffic.
Mild Green	Medium	Port is handling moderate traffic.
Light Green	Low	Port is handling light traffic.
Grey	Off	Port is unused.
White	No data	No data available for the port.

Use the heatmap to:

- Identify unused ports (grey), which are candidates for shutdown recommendations. See, ["Optimize Energy Consumption by Shutting Down Unused Ports" on page 285](#).
- Identify unused line cards (grey), which are candidates for shutdown recommendations. See, ["Optimize Energy Consumption by Shutting Down Unused Line Cards" on page 298](#).
- Use forecasted traffic patterns to spot recurring low-utilization windows and determine when specific ports can be shut down temporarily without impacting network traffic. See, ["Optimize Low Traffic Ports Based on Historical Traffic Patterns" on page 289](#).
- Identify ports with consistently low traffic (light green) that reduces device's energy efficiency and schedule port shutdown.
- Map per-port trends to power-throughput behavior over time

For example, if a port remains light green or gray across multiple hours, and the device's power usage does not drop during those periods, the port is contributing to idle power draw.

About the Power Supply Tab

The **Power Supply** tab provides insights into the health, status, and efficiency of the power modules installed on the device. These modules are typically hot-swappable and often deployed in redundant configurations (N+1) to ensure power availability.


Network devices usually support multiple power modules to maintain uptime even if one module fails. Each module may connect to one or more input power sources, such as AC mains or DC battery. The number of module slots can vary depending on the device model and usually can be replaced without powering off the device.

Use the tab to monitor power module performance and redundancy, diagnose input feed issues, and plan upgrades or maintenance based on module health.

- **View the Power Capacity of a Device**

The **Power Usage** widget uses a progress bar to display the device's maximum power capacity and its current power usage.

Figure 82: Power Usage Bar



- The full bar indicates the maximum power capacity of the device.
- **Green** segment represents the actual power usage of the device in Watts.
- **White** segment represents the remaining capacity in Watts.
- The percentage below the bar indicates how much of the maximum capacity is utilized by the device.

- The timestamp indicates when the data was last updated.

Use this view to understand the device's overall power capacity and current utilization by the device.

- **View Power Supply Module Availability**

The **Power Supply Modules** widget uses a donut chart to summarize the status of the power modules in the device.

For status definition, see [Table 71 on page 269](#).

Table 71: Power Module Status


Color	Status	Definition	Recommended Action
Green	Online	Power module is installed and supplying power.	No action required; verify redundancy
Grey	Offline	Power module is installed but not supplying power.	Check input feed, cabling, or module health

Use this chart to get a quick overview of power module availability and redundancy in the device.

In [Figure 8 on page 269](#), the entire ring is green, meaning 100% of modules are online.

If there were grey segments in the ring, it would indicate some modules are offline and may require troubleshooting.

Figure 83: Power Supply Module Chart




- **View Details on the Power Supply Modules**

The **Power Supply Modules** table provides detailed information on the individual power modules installed on the device.

The timestamp indicates when the data was last updated.

Use this table to assess module health, prevent power outages, and monitor module efficiency to optimize energy usage.



For field description, see [Table 72 on page 270](#)

Table 72: Power Module Table

Field	Description
Name	Name of the power module, along with an icon representing the module. Color coding is used to indicate the health of each module. • Green—Indicates the module is healthy. • Red—Indicates the module is faulty.
Power Load Bar	Visual indicator that shows the current power output of each module relative to its maximum capacity. The usage is measured in kilowatts (kW) and percentage, showing how much of the module’s capacity is being utilized. Color Scale • Green—Actual power being used by the module. • White—Remaining unused capacity. • Gray—Indicates the full bar but no active load. This is displayed for offline modules.

Table 72: Power Module Table *(Continued)*

Field	Description
Status	Indicates the status of the power module, whether the module is online or offline. See Table 71 on page 269 .
Efficiency	Shows the power conversion efficiency of each module, expressed in percentage. Higher values indicate more efficient module performance.

Use this table to ensure the device is not nearing power capacity thresholds, detect underperforming or degraded power modules, and plan for redundancy or upgrades.

For example, if a switch consumes 22 kW out of a total of 120 kW of power, it is operating at approximately 18% utilization. If the chassis provides two module slots but only one module is online, the device has no redundancy, and a single module failure would cause a complete outage. The dashboard helps you identify such risks and plan accordingly.

About the Thermal Profile Tab

The **Thermal Profile** tab displays information on the temperature, cooling performance, and sensor behavior of the device and how they contribute to energy usage.

Use this tab to monitor thermal performance, prevent overheating, and optimize fan and power usage of your device.

- **View a Device's Average Temperature and Average Fan Speed**

Use the **Device Avg Temperature and Avg Fan Speed** time-series graph to monitor trends in a device's temperature and cooling behavior.

Figure 84: Device Avg Temperature and Avg Fan Speed Graph




Table 73: Field Description

Axis	Metric	Description	Color	Unit
Left y-axis	Average Temperature	Shows the average temperature of the device over the selected time range. NOTE: Click the Avg Temperature legend to view only the temperature trend line on the graph. Double click the metric to clear the filter.	Blue	°C (Celsius)
Right y-axis	Average Fan Speed	Shows the average fan speed over the selected time range. NOTE: Click the Avg Fan Speed legend to view only the fan speed trend line on the graph. Double click the metric to clear the filter.	Brown	RPM
x-axis	Time Range (Today, Yesterday, Last 7 Days, Last 30 Days, Custom)	Represents the selected time range. By default, it is set to Last 30 Days.	—	Hours / Days



NOTE: Hover over any point on the lines to view the exact temperature and fan speed values for that timestamp.

Under normal operating conditions, the fan speed remains proportional to the device temperature.

Use this graph to identify:

- Periods where fan speed remains high while the temperature is low.
- Instances where temperature spikes are not met with a corresponding fan speed increase. This signals potential overheating risks or an issue with the thermal control system.
- Optimize fan speeds based on real-time needs to reduce the device's power draw and minimizes the heat load placed on the main data center cooling infrastructure.
- **View Power Trend of a Device**

The **Power Trend** line chart is a time-series graph that displays the power consumption of the device over a selected time period.

Figure 85: Power Trend Chart



The y-axis represents the power draw in Watts (W) and x-axis represents the selected time period (Today, Yesterday, Last 7 Days, Last 30 Days, or Custom).



TIP: Analyze the chart alongside the **Device Avg Temperature and Avg Fan Speed** chart to understand the relationship between power draw and thermal load of the device. Under normal operating conditions, the device's fans must respond with proportional increase in fan speed to prevent overheating and ensure thermal stability. If the thermal profile is stable but the power curve is consistently high, it suggests inefficient power to traffic proportionality and as a result increase the heat load.

- **Monitor the Top Sensors on the Device by Temperature Deviation**

The **Top Sensors by Temperature Deviation** widget provides a tabular view of the top five component sensors of the device that require immediate attention.

Table lists sensors with their current temperature and temperature deviation from baseline. You can also sort the table based on device temperature.

Figure 86: Top Sensors by Temperature Deviation Table

Top Sensors by Temperature Deviation		
Last updated December 02, 2025, 05:02:46 PM		
Sensor Name	Temperature Deviation	Temperature
FPC 0 et-0/0/31	0.00 °C	55.00 °C
FPC 0 et-0/0/29	-2.00 °C	53.00 °C
FPC 0 et-0/0/28	-3.00 °C	52.00 °C
FPC 0 et-0/0/30	-4.00 °C	51.00 °C
FPC 0 Sensor TD4 Max Reading	-4.50 °C	50.50 °C

Use this table to identify sensors that have the highest fluctuation or variation from a stable or baseline temperature.

For field description, see [Table 74 on page 274](#).

High deviation suggests potential hotspots and thermal issues, which could be caused by sudden spikes of workload or a failing fan. High deviation places significant stress on the component over time and impacts device performance.

If components are frequently overheated and show high deviation, it means that the device's cooling system (fans) or the facility's cooling (airflow) are failing to maintain stable operating conditions.

You can also correlate the sensor with the Power Trend Chart. A high-deviation sensor indicates that the power draw is likely spiking rapidly and frequently, generating inconsistent heat.

Table 74: Sensors Table Field Description

Field	Description
Sensor Name	Identifies the specific sensor being measured. This allows you to identify the source of the heat.

Table 74: Sensors Table Field Description (*Continued*)


Field	Description
Temperature Deviation	The variation between the device's measured temperature and the optimal temperature range. The icons help you quickly identify whether the device is running normal, warmer, or hotter than expected baseline. • Green—Indicates the device temperature is within expected range. No action required. • Yellow—The device is running slightly above its normal temperature but not yet critical. • Red—Temperature is significantly above normal and may impact device performance.
Temperature	The current aggregate temperature recorded by the sensor. Click on the arrowhead next to the header to sort the table in ascending or descending order.

- **View Fan Status of a Device**

The **Fan Status** widget shows the operational status of all fan trays in the device. The widget displays individual fans within each tray, their speed (RPM), and health indicators to help monitor thermal performance.

Use this widget to monitor the cooling performance of the fans. Failed fans reduce cooling efficiency of the device.

Figure 87: Fan Status View



You can sort the view by RPM or fan status to quickly identify the fans that are drawing the most power or those that have shut down completely due to low heat load.

Each tray is represented as a separate tile and inside each tray, individual fans are listed.

Revolutions Per Minute (RPM) indicates the rotational speed of each fan.

Color-coding is used to convey the health of each fan component. See, [Table 75 on page 276](#).

Table 75: Fan Health Indicator

Color	Status	Operational Status
Green	Ok	Fan is healthy and operating normally.
Yellow	Absent	Fan is not present in the tray.
Amber	Check	An error has been detected in the fan and may require inspection.
Red	Failed	The fan has failed or disconnected and requires immediate attention to prevent device from overheating.

Use the insights provided in this section to take action, whether it's replacing a component or adjusting the device's fan control policies. Thermal optimization can also help reduce cooling costs.

About the Recommendation Tab

The **Recommendation** tab provides a wide range of device-level energy optimization suggestions and their potential energy savings, based on telemetry collected from the device. Currently, you can get recommendations at the line card, PFE, and port levels.


Each recommendation is designed to streamline decision-making and give you the operational context you need to optimize energy usage.

The recommendation service executes once per day at around 3 AM UTC. Each execution uses historical port traffic telemetry and the configuration data available at the time of execution. The Recommendation tab displays the most recently generated results for the device.



NOTE: The service generates recommendations based on collected operational data. If adequate data is not yet available, the corresponding recommendations may not appear.

Figure 88: Recommendation Tab



At the top of the page, you'll see a set of recommendation cards, each representing a different category of energy optimization recommendation available for the device. Each card highlights the recommendation type, the component it applies to, and the potential energy savings you can achieve by applying this recommendation.

Click a card to open the detailed view panel for that recommendation type.

Each recommendation is presented in a structured format for easy evaluation. Use this information to assess and identify the most effective energy optimization actions for your network device.

 **NOTE:** The Recommendation tab appears empty, if no safe or meaningful recommendations or power saving actions were identified based on the latest device telemetry.

For field description, see [Table 76 on page 277](#).

Table 76: Recommendation Insights

Header	Description
Recommendation Card	Highlights the recommendation type, the count of device components that qualify for the recommendation, and the potential energy savings achievable if the recommendation is applied.
Device <i>Component</i>	Device component identified for power optimization. Use the drop-down to select a component and view its specific recommendation.

Table 76: Recommendation Insights (*Continued*)

Header	Description
Description	A summary of the observations that resulted in this recommendation.
Actions	Provides step-by-step guidance and relevant configlet-based actions required to apply the recommended changes on the device. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i>.
Benefits	Displays quantified advantages achievable by applying the action, such as estimated power savings, cost reduction, and reduced GHG emissions.
Risks	Highlight caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Indicates when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

Currently supported recommendations:

- ["Optimize Energy Consumption by Shutting Down Unused Ports" on page 285](#)
- ["Optimize Low Traffic Ports Based on Historical Traffic Patterns" on page 289](#)
- ["Optimize Low Traffic Ports Based on Forecasted Traffic Patterns" on page 293](#)
- ["Optimize Energy Consumption by Shutting Down Unused Line Cards" on page 298](#)
- ["Optimize Energy Consumption by Shutting Down Unused PFEs" on page 303](#)

Configure Port Threshold for the Organization

Administrators can use the **Port Threshold** setting to define the traffic utilization thresholds for device ports. Utilization is the percentage of traffic the port carries relative to its operating capacity. The

thresholds are set as percentage ranges and determine whether the port utilization is high, medium, or low.

Use the widget to define percentage ranges that classify port usage levels. By default, threshold for low traffic is set at 10% and high traffic at 80%.

The system uses these utilization ranges to classify port behavior in [Per Port Peak / Off-Peak Traffic Monitoring on page 266](#) heatmap and to generate current and forecasted low traffic port optimization recommendations.

For port utilization description, see [Table 77 on page 279](#).

Table 77: Port Utilization Category

Utilization	Description
Low	The port is handling very little traffic relative to its capacity. Ports consistently falling in this range may be idle or underutilized ports and qualifies for optimization recommendations.
Medium	The port is carrying a moderate amount of traffic. Ports in this range are active but not operating in full capacity.
High	The port is carrying high volumes of traffic relative to its capacity. Ports in this range are heavily used and are not candidates for optimization recommendations.

You can configure global port utilization thresholds from the **Organization Settings** page. The configuration automatically applies to all sites and all devices in the organization unless overridden by site-specific settings. The system provides default threshold values, inferred from historical traffic and usage, which you can adjust.




TIP: We recommend setting thresholds that reflect your organization's traffic patterns and operational behavior instead of relying on default values.

To configure port thresholds at the organization-level:

1. Navigate to **Organization > Organization Settings**.
2. Locate the **Port Threshold** slider.

Figure 89: Port Traffic Threshold



For description, see [Widget Components on page 280](#).

Table 78: Widget Components

Component	Description
Slider	Defines the boundaries between Low, Medium, and High utilization. If you do not configure port thresholds, default thresholds are applied for the ports.
Percentage markers	Indicate where each utilization range starts and ends on the 0–100% scale. NOTE: By default, the Low utilization range extends up to 10% , Medium from 11% up to 79% , and High from 80% up to 100% of the full port utilization scale.
Color bands on the slider	Visually represents each utilization range. The same colors appear in Per Port Peak / Off-Peak Traffic Monitoring heatmap.
Color legend	Shows how Low, Medium, and High utilization levels for a port are represented on the slider and across the Energy dashboards. • Light Green—Low utilization • Medium Green—Medium utilization • Dark Green—High utilization

3. Adjust the slider to set the utilization thresholds.

The slider has two draggable handles:

- **Left Handle**—Drag the handle to control the boundary between **Low** and **Medium** utilization.

Move this handle to the right to increase the Low utilization range (light green band expands).

Move this handle to the left to increase the Medium utilization range (medium green band expands).

- **Right Handle**—Drag the handle to control the boundary between **Medium** and **High** utilization.

Move this handle to the right to increase the Medium utilization range (medium green band expands).

Move this handle to the left to increase the High utilization range (dark green band expands).

The slider displays the percentage at each handle position, indicating how much of the scale each utilization range occupies.

For example, if the left handle is set at **30%** and the right handle at **70%**, the system interprets the ranges as:

- Ports with traffic **less than or equal to 30%** are considered low utilization ports.
- Ports with traffic **greater than 30% and up to 70%** are considered medium utilization ports.
- Ports with traffic **above 70%** are considered as high utilization ports.

4. Click **Save** to apply the changes.

Once applied, the dashboard updates the **Per Port Peak / Off-Peak Traffic Monitoring** heatmap of all devices to reflect current and forecasted utilization based on the configured thresholds.

Updating these thresholds may change how ports are classified and can result in updated "[Low Traffic Port Optimization](#)" on page 289 and "[Forecasted Low Traffic Port Optimization](#)" on page 293 recommendations for all devices in the organization, including changes to the time windows during which the recommendation is applicable.



REMEMBER: Update port thresholds as needed to reflect changes in your organization's traffic patterns or energy optimization goals.

RELATED DOCUMENTATION

[Understanding the Device-Level Sustainability Dashboard](#) | 258

Configure Port Threshold for a Site

Administrators can use the **Port Threshold** widget to define the traffic utilization thresholds for device ports. Utilization is the percentage of traffic the port carries. The thresholds are set as percentage ranges and determine whether the port utilization is high, medium, or low.

Use the widget to define percentage ranges that classify port usage levels. By default, threshold for low traffic is set at 10% and high traffic at 80%.

The system uses these utilization ranges to classify port behavior in Per Port-Peak/ Off-Peak Traffic heatmap and to generate current and forecasted low traffic port optimization recommendations.

For port utilization description, see [Table 79 on page 282](#).

Table 79: Port Utilization Category

Utilization	Description
Low	The port is handling very little traffic relative to its capacity. Ports consistently falling in this range may be idle or underutilized ports and qualifies for optimization recommendations.
Medium	The port is carrying a moderate amount of traffic. Ports in this range are active but not operating in full capacity.
High	The port is carrying high volumes of traffic relative to its capacity. Ports in this range are heavily used and are not candidates for optimization recommendations.

The system provides default thresholds at the organization-level, inferred from historical traffic and usage patterns. These thresholds automatically apply to all sites unless they are customized at the site-level. Site-level configurations override the global thresholds.




TIP: We recommend setting site-specific thresholds to reflect individual site's traffic patterns and operational behavior instead of using default values.

To configure port thresholds at the site-level:

1. Navigate to **Organization > Site Configuration**.
2. Click a site for which you want to set the port thresholds.
3. Locate the **Port Traffic Threshold** slider.

Figure 90: Port Threshold



For description, see [Widget Components on page 283](#).

Table 80: Widget Components

Component	Description
Slider	Defines the boundaries between Low, Medium, and High utilization. If you do not configure port thresholds, default thresholds are applied for the ports.
Percentage markers	Indicate where each utilization range starts and ends on the 0–100% scale. NOTE: By default, the Low utilization range extends up to 10% , Medium up to 79% , and High up to 100% of the full port utilization scale.
Color bands on the slider	Visually represents each utilization range. The same colors appear in the site's Per Port Peak / Off-Peak Traffic Monitoring heatmap.
Color legend	Shows how Low, Medium, and High utilization levels are represented across Energy dashboards. • Light Green—Low utilization • Medium Green—Medium utilization • Dark Green—High utilization

4. Adjust the slider to set the utilization thresholds.
- The slider has two draggable handles.

- **Left Handle**—Drag the handle to control the boundary between **Low** and **Medium** utilization.

Move this handle to the right to increase the Low utilization range (light green band expands).

Move this handle to the left to increase the Medium utilization range (medium green band expands).

- **Right Handle**—Drag the handle to control the boundary between **Medium** and **High** utilization.

Move this handle to the right to increase the Medium utilization range (medium green band expands).

Move this handle to the left to increase the High utilization range (dark green band expands).

The slider displays the percentage at each handle position, indicating how much of the scale each utilization range occupies.

For example, if the left handle is set at **30%** and the right handle at **70%**, the system interprets the ranges as:

- Ports with traffic **less than or equal to 30%** are considered low utilization ports.
- Ports with traffic **greater than 30% and up to 70%** are considered medium utilization ports.
- Ports with traffic **above 70%** are considered as high utilization ports.

5. Click **Save** to apply the configuration.

Once applied, the dashboard updates the **Per Port Peak / Off-Peak Traffic Monitoring** heatmap of device's in the site to reflect utilization based on the configured thresholds.

Updating these thresholds may change how ports are classified and can result in updated "[Low Traffic Port Optimization](#)" on page 289 and "[Forecasted Low Traffic Port Optimization](#)" on page 293 recommendations for devices in the site, including changes to the time windows during which the recommendation is applicable.



REMEMBER: Update port thresholds as needed to reflect changes in site's traffic patterns or energy optimization goals.

Optimize Energy Consumption by Shutting Down Unused Ports

SUMMARY

This topic describes how to optimize energy consumption by identifying and shutting down unused ports on network devices.

Network devices are typically provisioned with a large number of ports to ensure scalability and redundancy. However, many of these ports remain unused for extended periods but still continue to consume baseline power.

The **Device Port Power Optimization** recommendation list ports that show no signs of operational use and are safe to disable. By shutting down these ports, you can reduce the device's baseline energy consumption, lower operational costs, and reduce GHG emissions.

A port is selected for the **Device Port Power Optimization** recommendation when:

- It is administratively enabled (admin-state up) but has no active link (link-state down).
- A transceiver is inserted but the port does not carry any traffic.
- The operational state is up but the port shows no Tx/Rx activity.

The **Device Port Power Optimization** recommendation provides actionable guidance, such as the list of unused ports, benefits and risks associated with the optimization action, and the relevant configlet-based actions to safely disable ports and reduce idle power consumption. It also highlights the potential energy savings achievable by applying the recommendation.



NOTE:

The recommendation service runs once per day at around 3 AM UTC. Each run uses historical port traffic telemetry and configuration data available at the time. The Recommendation tab displays the most recently generated results for the device.

Changes to port states, traffic patterns, or threshold configurations that occur outside the current run window are reflected in recommendations after the next run.

Use the **Per-Port Peak / Off-Peak Traffic Patterns** heatmap available on the Device-Level Sustainability Dashboard to view traffic patterns for all ports.



NOTE: The Recommendation tab appears empty, if no safe or meaningful recommendations or power saving actions were identified based on the latest device telemetry.

To view unused ports and review the recommended actions:

1. From the main menu, click **Sustainability**.

The Organization-level Sustainability Dashboard appears.

2. Navigate to the **Top Sites by Energy Consumption** widget and click on a site with high energy consumption.

The Site-Level Sustainability Dashboard appears.

3. On the Site-Level Sustainability Dashboard, navigate to the **Top Device by Energy Consumption** widget and click on a device to investigate further.

The Device-Level Sustainability Dashboard appears.

4. On the Device-Level Sustainability Dashboard, review the **Per-Port Peak / Off-Peak Traffic Patterns** heatmap to understand the device's port activity trends.

5. Navigate to the **Recommendation** tab.

The **Device Port Power Optimization** recommendation card appears on the top of the page.

The card highlights the recommendation type, the count of ports identified as unused and safe to disable, and the potential energy savings you can achieve by applying this recommendation.

The recommendation identifies ports that are not actively carrying traffic. The potential savings indicate the estimated amount of energy that can be conserved by disabling unused ports that are not actively carrying traffic and by removing installed optical transceivers that are powered on but not in use.




NOTE:

- Recommendations are generated only if unused ports are identified on a device.
- To generate energy optimization recommendations for unused ports at least 24 hours of data is required.

Click the recommendation card to open the detailed recommendation panel.

[Figure 1 on page 287](#) shows a sample **Device Port Power Optimization** recommendation card and panel.

Figure 91: Sample Recommendation



The recommendation includes the following fields:

Table 81: Field Description

Field	Description
Recommendation Card	A high-level description of the recommendation type. Indicates the type of recommendation, the number of ports identified for applying the recommendation, and the potential energy savings achievable if the recommendation is applied.
Device Port	Port identified as unused and eligible for shutdown. Use the drop-down to select a port and view its specific recommendation.
Description	Summary of why the port is identified as unused and why the action is recommended, referencing energy and traffic data points.
Actions	Provides step-by-step guidance and relevant configlet-based actions required to disable the port without impacting the network traffic. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i>.
Benefits	Displays quantified advantages, such as estimated power and cost savings and reduced GHG emissions.

Table 81: Field Description (Continued)

Field	Description
Risks	Highlight caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Indicates when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

6. Apply the recommended action.

Use the configlets in Apstra Data Center Director to apply configuration changes to the device.

Configlets are custom configuration snippets that you can apply on a device through Apstra Data Center Director.



NOTE:

- Configuration changes made using configlets override Apstra-intended configuration.
- You may not get warnings for incorrect configurations pushed using configlets. Ensure you test and validate configlets in a dedicated validation environment before deploying them to the production environment.
- Passwords and other secret keys are not encrypted in configlets.

7. After applying the configuration changes, return to the Sustainability dashboards to validate the results.

Regularly verify device energy consumption, port activity, power savings, and so on.



NOTE: Changes to port states, traffic patterns, or threshold configurations may take up to 24 hours to be reflected in the Sustainability dashboards.

Optimize Low Traffic Ports Based on Historical Traffic Patterns

SUMMARY

This topic describes how to reduce device-level energy consumption by disabling ports that consistently carry low or negligible traffic during specific time windows, based on historical usage patterns.

Network traffic often exhibits predictable patterns over time. Some ports consistently carry low or negligible traffic during specific off-peak windows. These idle ports continue to stay powered on and consume baseline power.

The **Low Traffic Port Optimization** recommendation list such ports by analyzing several days of historical per-port traffic telemetry. When the system detects a stable pattern of repeated low utilization during defined time windows for a port, it marks the port as a safe candidate for temporary shutdown during those windows.

Disabling such ports helps reduce the device's baseline energy consumption without affecting workload performance.

By aligning each port's operational state with its actual traffic demand, this recommendation helps make the device's power usage more proportional to the network's utilization patterns, improve energy efficiency, and lower operational cost.

A port is selected for the **Low Traffic Port Optimization** recommendation when:

- It consistently operates below the configured low-traffic utilization threshold during the same hours across multiple consecutive days.



NOTE: You can configure the port utilization thresholds used for this analysis at the organization and site levels. See, "[Configure Port Threshold for the Organization](#)" on page 278 and "[Configure Port Threshold for a Site](#)" on page 282.

If no thresholds are configured, the system applies the default thresholds for all ports.

By default, the threshold for low traffic is set at 10%.

- Disabling the port during the low-traffic window is operationally safe.

- Sufficient historical telemetry exists to establish a predictable pattern.

The **Low Traffic Port Optimization** recommendation provides actionable guidance, such as the list of eligible ports, the observed low-traffic windows, and the relevant configlet-based actions required to safely disable ports and reduce idle power consumption. It also highlights the potential energy savings achievable by applying the recommendation.



NOTE: The recommendation service runs once per day at around 3 AM UTC. Each run uses historical per-port traffic telemetry and the configuration data available at the time. The Recommendation tab displays the most recently generated results for the device. Changes to port states, traffic patterns, or threshold configurations that occur outside the current run window are reflected in recommendations after the next run.

Use the **Per-Port Peak / Off-Peak Traffic Patterns** heatmap available on the Device-Level Sustainability Dashboard to view traffic patterns for all ports.



NOTE: The Recommendation tab appears empty, if no safe or meaningful recommendations or power saving actions were identified based on the latest device telemetry.

To view underutilized ports and review the recommended actions:

1. From the main menu, click **Sustainability**.
The Organization-level Sustainability Dashboard appears.
2. Navigate to the **Top Sites by Energy Consumption** widget and click on a site with high energy consumption.
The Site-Level Sustainability Dashboard appears.
3. On the Site-Level Sustainability Dashboard, navigate to the **Top Device by Energy Consumption** widget and click on a device to investigate further.
The Device-Level Sustainability Dashboard appears.
4. On the Device-Level Sustainability Dashboard, review the **Per-Port Peak / Off-Peak Traffic Patterns** heatmap on the **Overview** tab to understand the device's port activity trends.
5. Navigate to the **Recommendation** tab.
The **Low Traffic Port Optimization** recommendation card appears on the top of the page.

The card highlights the recommendation type, the count of ports identified as candidates for temporary shutdown, and the potential energy savings you can achieve by applying this recommendation. The potential savings indicate the estimated amount of energy that can be conserved by temporarily disabling network ports that exhibit consistently low traffic over time. This also includes the associated reduction in power consumption from optical transceivers linked to these ports.



NOTE:

- Recommendations are generated only if low traffic ports are identified on a device.
- Atleast 3 days of data is required to generate low traffic port optimization recommendation.

Click the recommendation card to open the detailed recommendation panel.

The recommendation includes the following fields:

Table 82: Field Description

Field	Description
Recommendation Type	A high-level description of the recommendation type. Indicates the type of recommendation, the number of ports identified for applying the suggestion, and the potential energy savings achievable if the recommendation is applied.
Device Port	Ports identified as carrying low or negligible traffic. Ports are classified into the following types: NOTE: The port type is shown only for connected ports. • Server—Ports that connect network devices to end hosts. • Fabric—Ports that interconnect switches within the data center fabric. • Unknown—Ports that do not belong to any Apstra-managed topology or ports for which topology or device-role information is not available. Use the drop-down to select a port and view its specific recommendation.
Summary	A brief explanation of the observed traffic patterns during specific time windows and the recommended optimization action.

Table 82: Field Description (*Continued*)

Field	Description
Actions	Provides step-by-step guidance and relevant configlet-based actions required to disable the port without impacting the network traffic. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i> .
Benefits	Displays quantified advantages, such as estimated power and cost savings and reduced GHG emissions.
Risks	Highlight caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Shows when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

6. Apply the suggested recommendation.

Use the configlets in Apstra Data Center Director to apply configuration changes to the device.

[Configlets](#) are custom configuration snippets that you can apply on a device through Apstra Data Center Director.



NOTE:

- Configuration changes made using configlets override Apstra-intended configuration.
- You may not get warnings for incorrect configurations pushed using configlets. Ensure you test and validate configlets in a dedicated validation environment before deploying them to the production environment.
- Passwords and other secret keys are not encrypted in configlets.

7. After applying the configurations, return to the Sustainability dashboard to validate the results.

Regularly verify device's port activity, power savings, energy consumption, and so on.



NOTE:

- Changes to port states, traffic patterns, or threshold configurations may take up to 24 hours to be reflected in the Sustainability dashboards.
- Updating port thresholds affects future recommendations and may also change the applicability window of recommendations after the next execution cycle.

Traffic patterns can change over time, so we recommend regularly returning to the Sustainability dashboards to monitor actual port utilization patterns, and re-enable ports when traffic demand increases or operational requirements change.

RELATED DOCUMENTATION

[Optimize Low Traffic Ports Based on Forecasted Traffic Patterns](#) | 293

Optimize Low Traffic Ports Based on Forecasted Traffic Patterns

SUMMARY

This topic describes how to reduce device-level energy consumption by using forecasted port usage patterns to identify and disable ports during recurring low-traffic windows.

Network traffic often exhibits predictable patterns over time. Some ports consistently carry low or negligible traffic during specific time windows. The recommendation system analyzes available forecasted port traffic data to detect ports that consistently carry low traffic during specific time windows and generate energy optimization recommendations.

A port is selected for the **Forecasted Low Traffic Port Shutdown** recommendation when:

- It consistently operates below the configured low-traffic utilization threshold during the same hours across multiple days in the forecast.



NOTE: You can configure the port utilization thresholds used for this analysis at the organization and site levels. See, "[Configure Port Threshold for the Organization](#)" on page 278 and "[Configure Port Threshold for a Site](#)" on page 282.

If no thresholds are configured, the system applies the default thresholds for all ports.

By default, the threshold for low traffic is set at 10%.

- Disabling the port during the low-traffic window is operationally safe.
- Sufficient forecasted data exists to establish a predictable pattern.



NOTE: The recommendation service runs once per day at around 3 AM UTC. Each run uses forecasted per-port traffic data and the configuration data available at the time. The Recommendation tab displays the most recently generated results for the device.

Changes to port states, traffic patterns, or threshold configurations that occur outside the current run window are reflected in recommendations after the next run.

The **Forecasted Low Traffic Port Shutdown** recommendation:

- Identifies ports that repeatedly show low traffic forecast during the same hours across multiple days.
- Estimates potential energy savings if these ports are disabled during those off-peak windows.
- Provides configlet-based actions to safely apply configuration changes on the device without affecting network traffic.
- Helps align device power usage with actual network demand, reducing baseline energy consumption and operational costs.

Use the **Per-Port Peak / Off-Peak Traffic Patterns** heatmap on the Device-Level Sustainability Dashboard to view forecasted traffic patterns for all ports.



NOTE: The Recommendation tab appears empty if no safe or meaningful recommendations or power saving actions were identified based on the latest device telemetry.

To view forecasted traffic patterns and review the recommended actions:

1. From the main menu, click **Sustainability**.

The Organization-level Sustainability Dashboard appears.

2. Navigate to the **Top Sites by Energy Consumption** widget and click on a site with high energy consumption.


The Site-Level Sustainability Dashboard appears.

3. On the Site-Level Sustainability Dashboard, navigate to the **Top Device by Energy Consumption** widget and click on a device to investigate further.

The Device-Level Sustainability Dashboard appears.

4. On the Device-Level Energy Dashboard, enable **Forecast** and review the forecasted traffic patterns for the next 24 hours in the [Per-Port Peak / Off-Peak Traffic Patterns](#) on page 295 heatmap on the **Overview** tab to understand predicted traffic trends.

Figure 92: Heatmap with Forecast Enabled



5. Navigate to the **Recommendation** tab. The **Forecasted Low Traffic Port Shutdown** recommendation card appears on the top of the page.

The card highlights the recommendation type, the count of ports identified as candidates for temporary shutdown, and the potential energy savings you can achieve by applying this recommendation.

The potential savings indicate the estimated amount of energy that can be conserved by proactively disabling ports that are predicted to remain underutilized based on forecasted traffic patterns. This also includes the associated reduction in power consumption from optical transceivers linked to these ports.




NOTE:

- Recommendations are generated only if eligible ports are detected.
- Forecast based recommendations may require up to 15 days of data after Edge onboarding or upgrade.

Click the recommendation card to open the detailed recommendation panel.

Figure 2 on page 296 shows a sample **Forecasted Low Traffic Port Shutdown** recommendation card and panel.

Figure 93: Sample Recommendation Card



The recommendation panel includes the following fields:

Table 83: Field Description

Field	Description
Recommendation Type	A high-level description of the recommendation type. Indicates the type of recommendation, the number of ports identified for applying the suggestion, and the potential energy savings achievable if the recommendation is applied.
Device Port	Port identified as carrying low or negligible traffic. Ports are classified into: NOTE: The port type is shown only for connected ports. • Server—Ports that connect network devices to end hosts. • Fabric—Ports that interconnect switches within the data center fabric. • Unknown—Ports that do not belong to any Apstra-managed topology or ports for which topology or device-role information is not available. Use the drop-down to select a port and view its specific recommendation.

Table 83: Field Description (*Continued*)

Field	Description
Summary	A brief explanation of the predicted traffic patterns during off peak time windows and the recommended optimization action.
Actions	Provides step-by-step guidance and relevant configlet-based actions required to disable the port without impacting the network traffic. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i> .
Benefits	Displays quantified advantages, such as estimated power and cost savings and reduced GHG emissions.
Risks	Highlight caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Shows when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

6. Apply the suggested recommendation.

Use the configlets in Apstra Data Center Director to apply configuration changes to the device.

[Configlets](#) are custom configuration snippets that you can apply on a device through Apstra Data Center Director.



NOTE:

- Configuration changes made using configlets override Apstra-intended configuration.
- You may not get warnings for incorrect configurations pushed using configlets. Ensure you test and validate configlets in a dedicated validation environment before deploying them to the production environment.
- Passwords and other secret keys are not encrypted in configlets.

7. After applying the configurations, return to the Sustainability dashboard to validate the results.

Regularly verify forecasted port activity, power savings, energy consumption and so on.

**NOTE:**

- Changes to port states, traffic patterns, or threshold configurations may take up to 24 hours to be reflected in the Sustainability dashboards.
- Updating port thresholds affects future recommendations and may also change the applicability window of recommendations after the next execution cycle.

Traffic patterns can change over time, so we recommend regularly returning to the Sustainability dashboards to monitor actual port utilization patterns and forecasts, and re-enable ports when traffic demand increases or operational requirements change.

RELATED DOCUMENTATION

[Optimize Low Traffic Ports Based on Historical Traffic Patterns](#) | 289

Optimize Energy Consumption by Shutting Down Unused Line Cards

SUMMARY

This topic describes how to optimize energy consumption by identifying and shutting down unused line cards on network devices.

Line cards (also referred to as Flexible PIC Concentrators (FPCs) on Juniper platforms) provide interface density and forwarding capacity to support scalability and redundancy. However, a chassis often contains multiple line cards that carry no traffic while continuing to consume significant power.

The recommendation service analyzes device telemetry to identify unused line cards that can be safely powered down and generate energy optimization recommendations for eligible line cards. You can view the recommendations on the **Device Linecard Power Optimization** recommendation card.



NOTE: The **Device Linecard Power Optimization** recommendation is available only for devices that support multiple linecards.

Currently, the recommendation is supported for the following platforms: PTX10016, PTX10008, PTX10004, PTX10001-36MR.

Disabling unused line cards provide one of the most impactful opportunities for reducing chassis-level power consumption.

A line card is selected for the **Device Linecard Power Optimization** recommendation when:

- All ports mapped on the line card show no traffic activity over a sustained observation window.
- No active production interfaces are mapped to the card.
- Redundancy policies and switchover capacity remain functional after removal.
- No critical services rely on the card.

The **Device Linecard Power Optimization** recommendation provides a list of unused line cards, historical utilization trends, configlet-based actions to safely take the line card offline, and benefits and risks associated with disabling the line card. It also highlights the potential energy savings achievable by applying the recommendation.



NOTE: The recommendation service runs once per day at around 3 PM UTC and uses the historical forwarding-plane telemetry and device configuration data available at that time. Changes in traffic demand or redundancy configuration are reflected after the next run.

By disabling unused line cards, you can align the active forwarding capacity with actual demand. This helps with significant reduction in total chassis power draw, lower fan speed and cooling overhead, improve overall power efficiency, and reduce operational cost.

To view unused line cards and review the recommended actions:

1. From the main menu, click **Sustainability**.

The Organization-Level Sustainability Dashboard appears.

2. Navigate to the **Top Sites by Energy Consumption** widget and click on a site with high energy consumption.

The Site-Level Sustainability Dashboard appears.

3. On the Site-Level Sustainability Dashboard, navigate to the **Top Device by Energy Consumption** widget and click on a device to investigate further.

The Device-Level Sustainability Dashboard appears.

4. On the Device-Level Sustainability Dashboard, review the **Per-Port Peak / Off-Peak Traffic Monitoring** heatmap to understand the device's port activity trends across line cards.

Figure 1 on page 300 shows the Per-Port Peak / Off-Peak Traffic heatmap for line card 0. Several cells appear green, indicating active traffic on those ports during the observed 24-hour window. Therefore, this line card is not eligible for shutdown.

Figure 94: Per-Port Peak / Off-Peak Traffic Heatmap for an Active Line Card





Figure 2 on page 300 shows the ports associated with line card 1. All cells appear grey, indicating no traffic activity during the observed 24-hour window. This makes the line card eligible for shutdown.

Figure 95: Per-Port Peak / Off-Peak Traffic Heatmap for an Unused Line Card



5. Navigate to the **Recommendation** tab.
- The **Device Linecard Power Optimization** recommendation card appears on top of the page.

The card highlights the recommendation type, the count of line cards identified as unused and safe to disable, and the potential energy savings you can achieve by applying this recommendation.




NOTE:

- Recommendations are generated only if unused line cards are identified on a device.
- To generate energy optimization recommendations for unused line cards, atleast 24 hours of data is required.

Click the recommendation card to open the detailed recommendation panel.

Figure 3 on page 301 shows a sample **Device Line Card Power Optimization** recommendation card and panel.

Figure 96: Device Line Card Power Optimization Sample Recommendation



The recommendation panel includes the following fields:

Table 84: Field Description

Field	Description
Recommendation Card	A high-level description of the recommendation type. Indicates the type of recommendation, the number of line cards identified for applying the recommendation, and the potential energy savings achievable if the recommendation is applied.

Table 84: Field Description (*Continued*)

Field	Description
Device FPC	Line cards identified as unused and eligible for shutdown. Select a line card from the drop-down list to view the recommended action.
Description	Summary of why the line card is identified as unused and why the action is recommended, referencing energy and traffic data points.
Actions	Provides step-by-step guidance and relevant configlet-based actions required to disable the line card without impacting the network traffic. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i> .
Benefits	Displays quantified advantages, such as estimated power and cost savings and reduction in GHG emissions.
Risks	Highlights caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Indicates when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

6. Apply the recommended action.

Use the configlets in Apstra Data Center Director to apply configuration changes to the device.

[Configlets](#) are custom configuration snippets that you can apply on a device through Apstra Data Center Director.



NOTE:

- Configuration changes made using configlets override Apstra-intended configuration.

- You may not get warnings for incorrect configurations pushed using configlets. Ensure you test and validate configlets in a dedicated validation environment before deploying them to the production environment.
- Passwords and other secret keys are not encrypted in configlets.

7. After applying the configuration changes, return to the Device-Level Sustainability Dashboard to validate the results, such as reduction in device-level energy consumption, power cost, and confirm traffic redistribution across remaining line cards.

Regularly review traffic patterns, assess capacity, ensure redundancy and failover policies.



NOTE: Changes to line card state and traffic patterns may take up to 24 hours to reflect on the Sustainability Dashboards.

For granular, per-port energy optimization, see ["Optimize Energy Consumption by Shutting Down Unused Ports" on page 285](#).

For forwarding-plane energy optimization, see ["Optimize Energy Consumption by Shutting Down Unused PFEs" on page 303](#).

Optimize Energy Consumption by Shutting Down Unused PFEs

SUMMARY

This topic describes how to optimize energy consumption by identifying and shutting down unused Packet Forwarding Engines (PFEs) on network devices.

Most network devices are designed with multiple Packet Forwarding Engines (PFEs) to support scalability and redundancy. However, the available forwarding capacity may not always be fully utilized. Unused PFEs continue to consume significant baseline power even when carrying little or no traffic.

The recommendation service analyzes device telemetry to identify unused PFEs that can be safely powered down and generates energy optimization recommendations for eligible PFEs. You can view the recommendations on the **Device PFE Power Optimization** recommendation card.

The recommendation identifies PFEs that are consistently unused and can be safely powered down without impacting traffic forwarding or redundancy.



NOTE: Currently, the **Device PFE Power Optimization** recommendation is supported on the following platforms: PTX10016, PTX 10008, PTX10004, PTX10001-36MR.

A PFE is selected for the **Device PFE Power Optimization** recommendation when the following conditions are met:

- All ports mapped to the PFE show no traffic activity over a sustained observation window.
- No ports associated with the PFE have an active physical link.
- No production interfaces are operationally dependent on the PFE.
- Redundancy requirements and failover capacity remain functional after disabling the PFE.

The **Device PFE Power Optimization** recommendation provides a list of unused PFEs, historical utilization trends, configlet-based actions to safely take the PFE offline, and benefits and risks associated with disabling the PFE. It also highlights the potential energy savings achievable by applying the recommendation.



NOTE: The recommendation service runs once per day at around 3 AM UTC and uses the historical forwarding-plane telemetry and device configuration data available at that time. Changes in traffic demand or redundancy configuration are reflected after the next run.

By disabling unused PFEs, you can align the active forwarding capacity with actual demand. This helps reduce forwarding plane power consumption, lower cooling overhead, reduce operational cost, and improve overall data center sustainability.

To view unused PFEs and review recommended actions:

1. From the main menu, click **Sustainability**.

The Organization-Level Sustainability Dashboard appears.

2. Navigate to the **Top Sites by Energy Consumption** widget and click on a site with high energy consumption.

The Site-Level Sustainability Dashboard appears.

3. On the Site-Level Sustainability Dashboard, navigate to the **Top Device by Energy Consumption** widget and click on a device to investigate further.

The Device-Level Sustainability Dashboard appears.

4. Navigate to the **Recommendation** tab.

The **Device PFE Power Optimization** recommendation card appears on the top of the page.

The card highlights the recommendation type, the count of PFEs identified as unused and safe to disable, and the potential energy savings you can achieve by applying this recommendation.




NOTE:

- Recommendations are generated only if unused PFEs are identified on a device.
- To generate energy optimization recommendations for unused PFEs, atleast 24 hours of data is required.

Click the recommendation card to open the detailed recommendation panel.

Figure 1 on page 305 shows a sample **Device PFE Power Optimization** recommendation panel.

Figure 97: Device PFE Power Optimization Sample Recommendation



The recommendation panel includes the following fields:

Table 85: Field Description

Field	Description
Recommendation Card	A high-level description of the recommendation type. Indicates the type of recommendation, the number of PFEs identified for applying the recommendation, and the potential energy savings achievable if the recommendation is applied.

Table 85: Field Description (*Continued*)

Field	Description
Device PFE	PFE identified as unused and eligible for shutdown. Select a PFE from the drop-down list to view the recommended action. The drop-down lists entries in the format: FPC <number>/PFE <number>.
Description	Summary of why the PFE is identified as unused and why the action is recommended, referencing energy and traffic data points.
Actions	Provides step-by-step guidance and relevant configlet-based actions required to take the PFE offline without impacting the network traffic. NOTE: You must use the configlets in Apstra Data Center Director to apply configuration changes to the device. For more information on configlets, see Configlets in the <i>Apstra Data Center Director User Guide</i>.
Benefits	Displays quantified advantages, such as estimated power and cost savings and reduction in GHG emissions.
Risks	Highlights caveats such as potential dependencies, safety checks, and guardrails to consider before applying the suggested changes.
Timestamp	Indicates when the recommendation was generated and how long it is valid. NOTE: Recommendations are refreshed daily based on the latest device telemetry.

5. Apply the recommended action.

Use the configlets in Apstra Data Center Director to apply configuration changes to the device.

[Configlets](#) are custom configuration snippets that you can apply on a device through Apstra Data Center Director.



NOTE:

- Configuration changes made using configlets override Apstra-intended configuration.

- You may not get warnings for incorrect configurations pushed using configlets. Ensure you test and validate configlets in a dedicated validation environment before deploying them to the production environment.
- Passwords and other secret keys are not encrypted in configlets.

6. After applying the configuration changes to shutdown the PFE, return to the Device-Level Sustainability Dashboard to validate the results, such as reduction in device-level energy consumption and power cost.

Regularly review forwarding utilization trends, assess capacity, and ensure redundancy requirements remain satisfied after topology changes.



NOTE: Changes to PFE states and traffic patterns may take up to 24 hours to reflect in the Sustainability Dashboards.

For chassis-level optimization recommendation, see ["Optimize Energy Consumption by Shutting Down Unused Line Cards" on page 298](#).

For more granular, per-port energy optimization recommendation, see ["Optimize Energy Consumption by Shutting Down Unused Ports" on page 285](#).

Network Sustainability Analytics — Troubleshooting

IN THIS SECTION

- Problem: Sustainability Data Unavailable | 308
- Problem: Potential Energy Savings Data Unavailable | 308
- Problem: Power Cost Displayed as 0 | 309
- Problem: Sustainability Widget Unavailable on the Dashboard | 310
- Problem: Sites Not Displayed in Map View | 310
- Problem: Sustainability Data For Extended Periods Is Not Available | 311
- Problem: Data for Certain Months Unavailable in Sustainability Report | 311
- Problem: Incorrect Data Displayed When Selecting **Today** | 311
- Problem: Recommendations Unavailable After Enabling Sustainability Settings | 312
- Problem: Line Card and PFE Shutdown Recommendations Unavailable | 312

- Problem: Most Ports Shown in Light Green or Gray | 313
- Problem: Multiple Currency Support Issue | 313
- Problem: Port Type Differentiation Not Accurate (Fabric vs Server) | 314
- Problem: Efficiency of PEM Shown as N/A | 314

This topic helps you troubleshoot issues related to availability of sustainability data.

To view sustainability data, you must enable the Network Sustainability Analytics feature for your organization and configure the required settings. After the feature is enabled, it may take up to 24 hours for sustainability data to be available on the dashboards. For assistance, contact your designated HPE Networking representative.

Problem: Sustainability Data Unavailable

Description

Dashboard does not display sustainability data.

Solution

1. Verify the Apstra Edge version.

Sustainability features are supported only with Apstra Edge version 0.45.5 or higher.

2. If Apstra Edge is just onboarded and the status is displayed as **Registered** and **Connected**, wait for some time.

It might take over an hour for sustainability data to be available on the dashboard. For more information, see ["Verify DC Assurance Setup" on page 128](#).

Problem: Potential Energy Savings Data Unavailable

Description

The **Potential Energy Savings** widget shows no data across the organization, sites, and device Sustainability dashboards.

Solution

The projections on the Potential Energy Savings widget depend on the configured Average Power Cost and Carbon Intensity values, and the energy optimization recommendations.

Verify that the following values are set:

- Configure the power cost per unit either on the Organization Settings page or on the Site Configuration page.
- Modify the Carbon Intensity (gCO₂e/kWh) value as required either on the Organization Settings page or on the Site Configuration page. If Carbon Intensity is not configured, the default value of 443 will be used.

For more information, see ["Configure Sustainability Settings for the Organization" on page 213](#), ["Configure Sustainability Settings for a Site" on page 227](#).

The recommendations used to calculate potential savings require sufficient data before results begin to appear.

- Unused component (ports, line cards, PFEs) recommendations require at least 24 hours of data.
- Low usage port recommendations require up to 3 days of data.
- Forecast based recommendations may require up to 15 days of data after Edge onboarding or upgrade to generate relevant recommendations and potential savings data.

Once the required inputs are provided and sufficient data is available, the system can generate relevant optimization recommendations and provide meaningful potential savings insights.

Problem: Power Cost Displayed as 0

Description

Power cost is displayed as 0 on the Sustainability dashboards.

Solution

Configure the power cost per unit either on the Organization Settings page or on the Site Configuration page. For more information, see ["Configure Sustainability Settings for the Organization" on page 213](#), ["Configure Sustainability Settings for a Site" on page 227](#).

Problem: Sustainability Widget Unavailable on the Dashboard


Description

The DCA Dashboard does not display the Sustainability widget.

Solution

Click the **Edit** icon and select Sustainability from the list of widgets to be displayed on the Dashboard as shown in [Figure 98 on page 310](#).

Figure 98: Edit Dashboard Widgets



Problem: Sites Not Displayed in Map View

Description

Sites are not displayed in the map view of the Organization-Level Sustainability Dashboard.

Solution

1. Click **Organization > Site Configuration** and select the site.
2. In the **Location** tile, enter the latitude and longitude for the site.

For more information, ["Edit Site Configuration" on page 31](#).

Problem: Sustainability Data For Extended Periods Is Not Available

Description

Sustainability data for extended periods (daily, weekly, or monthly patterns) is not visible or appears incomplete.

Cause

Data availability for recently onboarded or upgraded Apstra Edge may be delayed, even when the device status shows **Registered** and **Connected**.

Data may also be unavailable due to connectivity issues across different scopes.

- At the organization level, insights may be missing if Apstra Edge is unavailable or if the available Apstra blueprint is incorrect.
- At the site level, data gaps may result from Edge or cloud connectivity issues.
- At the device level, insights may be missing if device status is not reflected correctly on the DC Edge page or if the device is listed as not connected.

Sustainability data is generated from historical data, and availability depends on the system completing the required data aggregation over selected time windows.

Problem: Data for Certain Months Unavailable in Sustainability Report

Description

The Sustainability report does not display data for certain months.

Cause

The Sustainability report shows energy data from the date on which sustainability feature is enabled for your organization. For example, if sustainability feature is enabled on July 15, the Sustainability report displays historical and forecasted data from July 15 to December 31 for that year. For more information, see ["Review the Sustainability Report" on page 251](#).

Problem: Incorrect Data Displayed When Selecting Today

The Sustainability dashboards may display inaccurate data if **Today** is selected in the time range selector.

Description

When you select **Today** from the time range selector, the data displayed may include extra hours or show fewer hours than expected.

Cause

The system uses **UTC** to determine the current day, while you may be located in a different time zone. This time zone difference causes the system's definition of **Today** (based on UTC) to not align with the your local calendar day. As a result, you may see extra hours of data or fewer hours of data, depending on your local time zone.

Problem: Recommendations Unavailable After Enabling Sustainability Settings

Description

Recommendations are not displayed immediately after enabling the Network Sustainability Analytics feature.

Cause

Recommendations are generated after periodic data collection. If sufficient data is not yet available, certain recommendations may not appear.

After enabling the Network Sustainability Analytics feature (Day 0):

- Unused port recommendations require atleast 24 hours of data.
- Low traffic port recommendations require one week of data and will be generated after seven days.
- Forecast based recommendations require atleast 14 days of data.

Recommendations are generated daily at around 3:00 AM UTC and continue to improve as more data becomes available. For more information, see ["About the Recommendation Tab" on page 276](#).

Problem: Line Card and PFE Shutdown Recommendations Unavailable

Description

Line card and PFE energy optimization recommendations are unavailable.

Solution

The line card and PFE recommendations are supported on platforms with multiple line cards. Currently, the recommendations are supported on the following platforms: PTX10016, PTX10008, PTX10004, PTX10001-36MR.

The recommendation service runs once per day and uses the historical forwarding-plane telemetry and device configuration data available at that time. Changes in traffic demand or redundancy configuration are reflected after the next run.

Problem: Most Ports Shown in Light Green or Gray

Description

In **Per Port - Peak / Off-Peak Traffic Monitoring**, most of the ports are shown in light green or gray. No port is highlighted in dark green.

Solution

By default, the threshold for low traffic is set at 10% and high traffic at 80%. Depending on your network requirements, you can configure traffic threshold at the site level or at the organization level from the **Organization Settings** or **Site Configuration** pages, respectively. For more information, see ["Configure Port Threshold for the Organization" on page 278](#), ["Configure Port Threshold for a Site" on page 282](#).

Problem: Multiple Currency Support Issue

Description

The Currency setting is not displayed at the site-level or appears as USD instead of the expected currency.

Solution

- The currency configured at the organization-level is applied to all sites in the organization. Ensure that the configured currency aligns with your regional, compliance, or financial reporting requirements.
- If the currency is not explicitly configured, the system defaults to USD.

Problem: Port Type Differentiation Not Accurate (Fabric vs Server)

Description

Port differentiation (Fabric vs Server) data is incorrect, missing, or not displayed for the energy optimization recommendation.

Cause

Port type differentiation is shown on the recommendation panel only for ports that are connected.

The port classification data is derived from Apstra blueprint.

The port type (Fabric vs Server) may not be shown on a recommendation if the port is not connected, if the blueprint data is incorrect or incomplete, or if there are connectivity issues with Apstra Edge.

Problem: Efficiency of PEM Shown as N/A

Description

In the Power Supply Modules widget, **Efficiency** is shown as N/A for PEM.

Cause

Power supplies of type PEM do not provide input power data. Hence, efficiency cannot be calculated for PEMs.

RELATED DOCUMENTATION

[Network Sustainability Analytics Overview | 201](#)

[Forecast Power Cost and GHG Emission | 203](#)

16

CHAPTER

DC Hub

IN THIS CHAPTER

- [DC Hub Overview | 316](#)
 - [About the DC Hub Page | 321](#)
-

DC Hub Overview

SUMMARY

This page describes how the DC Hub provides a comprehensive platform for the network administrator to discover, install, and manage add-ons to enhance Apstra functionalities.

IN THIS SECTION

- [DC Hub Overview | 316](#)
- [View Add-on Details | 317](#)
- [Import an Add-on | 319](#)
- [View My Activity | 320](#)
- [Benefits of DC Hub | 321](#)

DC Hub Overview

The DC Hub is a centralized platform within Juniper Data Center Assurance that helps you discover, access, and manage add-ons and enhancements for Apstra. These add-ons act as plug-and-play options to add functionality and analytics within Apstra. You can easily browse and import various add-ons directly into your AOS environment.


The additional configurations that can be deployed in Apstra Data Center Director are available across other repositories such as GitHub. The DC Hub acts as a single platform that provides you with consolidated add-ons containing the configurations required for specific functionalities. The DC Hub enables you to browse, import, and manage these add-ons directly onto your Apstra environment through a standardized process via the DC Assurance UI.

The DC Hub provides detailed descriptions, compatibility information, version details, and components for each add-on, allowing you to make informed decisions before installation. This streamlined import process, combined with post-installation guidance, ensures a seamless integration of add-ons into your existing Apstra environment.

The DC Hub currently hosts two types of add-ons:

- **AnalyticsPacks** — modules that include specialized dashboards, custom collectors, and probes to provide deeper network insights and monitoring capabilities
- **FeaturePacks** — modules that include configuration elements like Configlets and Property Sets that extend existing Apstra features for specialized requirements

Figure 99: DC Hub Analytics Pack



On the Hub, you can select either the **Analytics Pack** or the **Feature Pack** tabs to browse through the available add-ons. You can also use the **Search** bar to find specific add-ons as per your requirements.

You can toggle the display button on the top right of the page to view the available add-ons in list or table format. This page displays the name, version, description, and import status about each add-on. You can also select an add-on to view more details and to import it onto Apstra. The **My Activity** button takes you to a list of recorded import and unimport actions performed by you on the Hub.



NOTE: The add-ons that are developed and published by Juniper are tagged as Juniper Official.

View Add-on Details

Click the **View Details** link or icon to open the add-on page that provides information such as a brief description of the add-on, compatibility information, publisher's notes, and so on. You can also view a list of Edges where the add-on has already been imported.

Figure 100: Add-on Details Page




Table 86: Options on the Add-on Details Page.

Option	Description
1	Name of the add-on.
2	Select the version of the add-on.
3	Initiate the import process for the add-on.
4	Download the add-on to your local device's hard drive.
5	Open the repository where the add-on is hosted.

Table 86: Options on the Add-on Details Page. *(Continued)*

Option	Description
6	Go to My Activity page.
7	Juniper official – a tag that indicates that the add-on is developed by Juniper.
8	Date when the selected version of the add-on was published.
9	Description of the add-on.
10	List of Apstra versions that are supported for the add-on.
11	Apstra features that will be enhanced on importing the add-on.
12	Publisher notes with information about the add-on.
13	List of Apstra Edges where the add-on is already imported. Click the trash can icon to unimport the add-on from an instance.

Import an Add-on



NOTE: You must be a superuser to import add-ons from the Hub.

To import an add-on onto an Edge or a blueprint, perform the following steps:

1. Select an add-on from the Hub page and click **View Details**.
2. On the *<add-on>* page, click **Import**.
3. The **Import Add-on** dialog box opens with a list of supported Apstra versions, connected Edges, and their blueprints.

Figure 101: Import Add-on

Import DDoS

Choose an edge to select the blueprint into which you want to import.

Supported Apstra versions

5.0.0, 5.0.1, 5.1.0, 6.0.0

Filter

1-5 of 5

Edge Name	Resource Utilization	Apstra Version	Pack Compatibility	Edge Status
☒ Dogfood Apstra aos-2	Low	6.0.0	Yes	Connected
☐ SE_Demo_Lab	Low	6.0.0	Yes	Connected
☐ DFW-Demo	Low	5.0.0	Yes	Connected
☐ Dogfood Apstra	Low	6.0.0	Yes	Connected
☐ HOU-Demo	Low	5.1.0	Yes	Connected

Next

Cancel

4. Select the Edge where you want to import the add-on from the list and click **Next**.
5. Enter the appropriate details in the publisher-required input parameter fields. These parameters might differ for each add-on.

NOTE: You can import multiple instances of the same add-on onto a single Edge/blueprint with different input parameters for different requirements. But the identifiers must be unique for each instance of the import. If the identifiers for an import instances is the same as an existing instance, the latest import attempt will fail with an error.

6. Click **Import**.

NOTE: For Feature Packs, the imported add-on is found in the **Uncommitted** tab of the selected blueprint on Juniper Apstra. As a precaution, you must manually commit the imported changes for the add-on components to be integrated onto the Apstra instance. For more information, see [Commit / Revert Changes to Blueprint](#) in the *Apstra Data Center Director User Guide*.

View My Activity

You can view records of your import and unimport actions on the Hub. Click the **My Activities** tab from the Hub homepage or the **My Activities** icon on the top right of the page. A list of all the add-ons you have imported or unimported is displayed along with the add-on versions, blueprints, Edges, and status. You can also view the activity logs for each add-on and download the same in txt format.

Figure 102: My Activity Page

Pack Name	Pack Type	Pack Version	Publisher Name	Instance Name	Blueprint	Edge Name	Status	Imported Logs Time
DDoS	AnalyticsPack	1.0.0	Juniper	DDOS-2	Houston-demo	HOU-Demo	Success	Sep 17, 2025 6:05:58 PM
DDoS	AnalyticsPack	1.0.0	Juniper	DDOS-1	alpha40	Dogfood Apstra aos-2	Success	Sep 17, 2025 6:05:26 PM
duplicate-MAC-detection-timer	FeaturePack	1.0.0	Juniper		Dallas-Demo	DFW-Demo	Unimport Success	Aug 29, 2025 7:57:45 AM
DDoS	AnalyticsPack	1.0.0	Juniper	DDOS	Dallas-Demo	DFW-Demo	Unimport Success	Aug 29, 2025 7:56:59 AM
duplicate-MAC-detection-timer	FeaturePack	1.0.0	Juniper		Dallas-Demo	DFW-Demo	Unimport Success	Aug 29, 2025 7:28:05 AM
system-alarm	AnalyticsPack	1.0.0	Juniper	System_Alarm	Dallas-Demo	DFW-Demo	Unimport Success	Aug 29, 2025 7:27:56 AM
control-plane-DDoS-protection	FeaturePack	1.0.0	Juniper	Control Plane DDOS Pr...	Dallas-Demo	DFW-Demo	Unimport Success	Aug 29, 2025 7:27:44 AM
system-uptime	AnalyticsPack	1.0.0	Juniper	System_Uptime_test	Dallas-Demo	DFW-Demo	Unimport Success	Aug 28, 2025 8:09:27 PM

You can also select an add-on and unimport the instance from Apstra from the **My Activity** page.

Benefits of DC Hub

DC Hub provides the following benefits:

- Acts as a single, consolidated platform for add-ons that enhance existing Apstra features.
- Provides detailed information about available add-ons such as supported versions, components, Apstra version compatibility, and import statuses.
- Allows network administrators with a streamlined process to browse, import, manage, or unimport add-ons that improve Apstra functionality.

About the DC Hub Page

The Hub provides a centralized platform for discovering and managing add-ons that extend Apstra functionality, allowing network administrators to customize their deployments for specific requirements.

To access this page, click **Hub** on the main menu.

Tasks You Can Perform

- **Browse available add-ons** — Discover AnalyticsPacks and FeaturePacks that enhance Apstra functionalities.
- **Search and filter content** — Use sorting, filtering, and search functions to discover available add-ons that might be useful to you.
- **Import add-ons** — Deploy add-ons through your Apstra Edge to configure your Apstra environment with custom input parameters as per requirement.
- **Manage imported add-ons** — Track versions, perform upgrades, and reimport add-ons to access latest improvements and maintain compatibility with the Apstra version.
- **Download add-ons** — Save add-on files to your local device's hard drive for offline review or backup purposes.
- **Monitor Hub activities** — View your import and unimport history to track changes made to your environment.
- **Access activity logs** — Review and download activity logs for each add-on operation for troubleshooting and audit purposes.

17

CHAPTER

Service Level Expectations

IN THIS CHAPTER

- [Service Level Expectations Overview | 324](#)
 - [About the Service Levels Page | 333](#)
-

Service Level Expectations Overview

IN THIS SECTION

- [System Health | 326](#)
- [Link Health | 327](#)
- [Fabric Health | 328](#)
- [Virtual Infra Health | 329](#)
- [Service Health | 330](#)
- [Analyze SLE Scores | 330](#)
- [Probes to be Enabled in Apstra to View SLE Information | 331](#)

Service level expectations (SLE) define a benchmark for the performance of a network. An SLE consists of a set of attributes that provide information about the performance of a network. These attributes are known as classifiers. SLEs help network administrators understand whether the network is performing optimally and enable them to respond proactively to network events or performance issues.

Each classifier is assigned a value and the sum of the values for all the classifiers is the SLE score. The SLE score of an organization, a site, or a device is the aggregate score derived from all the classifiers defined for that SLE. The value or the percentage of the SLE score indicates the duration in minutes when the SLE meets or fails to meet the service levels. The SLE score is an indicator of how the network is performing. Administrators can analyze the SLE scores and look into events or problems that are impacting end user experience.

You can view SLE metrics at the organization level, site level, or at the device level.


To view information about SLE, click **Monitor > Service Levels**.

Juniper Data Center Assurance currently supports the following SLEs:

- System Health
- Link Health
- Fabric Health
- Virtual Infra Health
- Service Health

These SLEs measure the time when the SLE classifiers meet or do not meet the defined threshold, resulting in an improved user experience or a deterioration in user experience. Whenever there is an impact on user experience, administrators can use these metrics to proactively intervene and resolve issues. By observing the metrics provided by each classifier, you can easily conclude whether the network is performing optimally.

Figure 103: Service Levels Page



For each SLE, you can view

- **SLE Score**—SLE score for the organization, site, or the device during the selected period.
- Click **Success Rate** to view the SLE score in percentage, which is the percentage of time during which the user experience met the defined service level threshold. If the success rate is less than 100 percent, it indicates failures in a site or a device. For example, a service level of 99 percent indicates that the metric met the SLE goal 99 percent of the time and failed to meet the threshold 1 percent of the time.
- Click **Value** to view the value of the metric.



NOTE: The severity of an SLE metric can be in the range of 0.0 to 100.0, with 0.0 being the least severe and 100.0 being the most severe.


- **Timeline**—The timeline graph indicates how the network performed during the selected time range. Select a classifier or a sub-classifier to view the performance of that classifier, plotted on a graph. You can click and drag to select a specific time range on the graph or select a different time range from the drop-down to view more detailed information about the SLE. Mouse over the graph to view information about the overall service level, time, SLE score for the classifiers, and so on, displayed in a pop-up.
- **Classifiers**—Classifiers enable network administrators to perform a root cause analysis of the unsuccessful user experiences. A lower score for any of the classifiers can alert the administrators to address potential issues in the network.

The following sections describe SLEs and the classifiers that contribute to the SLE score for each SLE:

System Health

To access System Health SLE, click **Monitor > Service Levels > System Health**.

Figure 104: Classifiers and Sub-classifiers



System Health SLE provides information about how the device and its components are performing. Factors that impact the system health SLE include configuration changes, traffic passing through the device, performance of the device components such as the fans and power modules, and the system resource utilization.


- **Device Traffic**—Indicates that the device traffic affected the SLE score, leading to an impact on performance.

- Config Deviation—Change in system configuration affecting the SLE score.
- Environment—Problems with the system components affecting the SLE score. Has the following sub-classifiers:
 - Fan—Problems with the device's fan affecting the SLE score.
 - Power—Issues with the power module in the device affecting the SLE score.
 - Temp—Problems with system temperature affecting the SLE score.
- Resources—System resources that affect the SLE score. Has the following sub-classifiers:
 - CPU—Problems with CPU utilization affecting the SLE score.
 - Disk—Problems with disk utilization impacting the SLE score.
 - Memory—Problems with memory utilization affecting the SLE score.

Link Health

To access Link Health SLE, click **Monitor > Service Levels > Link Health**. Link Health SLE provides information about the interface states or connectivity issues that leads to a lower SLE score.

Figure 105: Link Health SLE




- Down Interfaces—An interface that is down impacts the SLE score.
- Bad Optics—Optics issues impact the SLE score.
- Hot Cold Interfaces—Hot cold interfaces impact the SLE score.
- Fabric Interfaces—Issues in fabric interfaces impact the SLE score.

- Specific Interfaces—Issues in specific interfaces impact the SLE score.
- Interface Flapping—Frequent change of state of an interface contributed to a lower SLE score
 - Fabric Interfaces—Issues in fabric interfaces impact the SLE score.
 - Specific Interfaces—Issues in specific interfaces impact the SLE score.
- LAG Issues—Issues in traffic flow on the link due to a link going down or due to excessive traffic on one of the links. A change in the classifier value indicates that the LAG Issues value has crossed the threshold configured in the probe in Apstra. It also shows which link is causing problems for traffic flow.
- Imbalance—A change in classifier value indicates that there is an imbalance in LAG usage.
- ESI—Issues in the ESI interfaces. A change in the classifier value indicates that the ESI value has crossed the threshold configured in the probe.

Fabric Health

To access Fabric Health SLE, click **Monitor > Service Levels > Fabric Health**. Fabric Health SLE provides information about the performance of the data center fabric, which includes the state of sessions between the leaf, spines, and superspines in the data center.




- BGP
 - IP Reachability Spine to Leaf—Problems with spine-to-leaf connectivity. A change in classifier value indicates there is a spine-to-leaf connectivity issue.
 - Sessions—Problems with BGP sessions either between a spine and leaf in the data center or with an external device, such as a superspine. A decrease in the classifier value indicates that the number of flaps has crossed the threshold value.
- EVPN—Reports anomalies related to EVPN routes in the network.

- **EVPN Type 2**—Problems in EVPN Type-2 routes. A decrease in the classifier value indicates that the expected EVPN Type-2 routes are missing on the device(s) in the network.
- **EVPN Type 3**—Problems in EVPN Type-3 routes on the leaf devices in the network. A decrease in the classifier value indicates that the expected EVPN Type-3 routes are missing on the device(s) in the network.
- **EVPN Type 5**—Problems with EVPN Type-5 routes on the leaf devices in the network. A change in the classifier value indicates that the expected EVPN Type-5 routes are missing on the device(s) in the network.
- **Flood List**—Problems with the flood list on the leaf devices in the network. A decrease in the classifier value indicates that the expected flood list forwarding entries are missing on the device(s) in the network.
- **ECMP**
 - **Imbalance**—Problems in selection of ECMP routes. A change in classifier value indicates there is an imbalance in usage of ECMP routes.

Virtual Infra Health

To access Virtual Infra Health SLE, click **Monitor > Service Levels > Virtual Infra Health**. Virtual Infra Health SLE provides information about deviation from intended behavior in the virtual infrastructure. To obtain Virtual Infra Health SLE metrics, you must configure vCenter when you adopt Apstra Edge.

Figure 106: Virtual Infra Health SLE




- **Config Mismatch**—Checks for mismatch between the Apstra-configured virtual networks in the data center fabric and the VLANs needed by the VMs running on the hypervisors attached to devices in the fabric. A decrease in the classifier value indicates that there are problems in VLAN configuration.
- **Non-redundant hosts**—Checks whether hosts have redundant hypervisors. A decrease in the classifier value indicates that not all hosts have redundant hypervisors configured.

Service Health

To access Service Health SLE, click **Monitor > Service Levels > Service Health**. Service Health SLE monitors and provides information about the overall service performance impact due to system health and link health related anomalies.

Figure 107: Service Health SLE



- **Impacted by System Issues**—Indicates that the system health related anomalies affected the SLE score. It uses the same metrics as the System Health SLE to calculate the impact on the services.
- **Impacted by Link Issues**—Indicates that the link health related anomalies affected the SLE score. It uses the same metrics as the Link Health SLE to calculate the impact on the services.

Analyze SLE Scores

You can analyze the SLE score from the Root Cause Analysis page. The Root Cause Analysis page provides visualizations for distribution, timeline, and statistics for service level failures and enables administrators to understand the impact of these issues. To view the Root Cause Analysis page, click **Monitor > Service Levels**. Then click an SLE or a classifier for more detailed information.

- **Statistics**—The Statistics tab displays the success rate of the SLE metric. Administrators can also view the distribution graph to understand the severity of the SLE with its impact duration.

- **Timeline**—The Timeline tab provides a graph that plots the SLE trend during the selected time range. Select a classifier or a sub-classifier to view the timeline graph. Mouse over the graph to view more detailed information about the SLE score.
- **Distribution**—Provides information about how a classifier impacts other devices in the data center.
- **Affected Items**—Provides information about the services, virtual networks, and routing zones.
 - **Interfaces**—Provides the list of interfaces that are impacted due to anomalies. This metric is available only for Link Health SLE.
 - **Services**—Provides the list of devices or services that failed to meet the defined service levels and also the impact of a specific device or service.
 - **Virtual Networks**—Provides the list of virtual networks that are impacted due to anomalies.
 - **Routing Zones**—Displays the list of routing zones that are impacted due to anomalies.

Probes to be Enabled in Apstra to View SLE Information

Table 87 on page 331 lists the predefined probes to be enabled in Apstra Data Center Director. For more information, see [Instantiate Predefined Probe](#) in the *Apstra Data Center Director User Guide*.

Table 87: Probes to be Enabled in Apstra

SLE	Probe	Affected Items
System Health	• Device System Health Probe • Drain Traffic Anomaly Probe • Device Environmental Checks	• Services • Virtual Networks • Routing Zones

Table 87: Probes to be Enabled in Apstra *(Continued)*

SLE	Probe	Affected Items
Link Health	• Optical Transceivers • Interface Flapping (Fabric Interfaces) • Interface Flapping (Specific Interfaces) • Hot/Cold Interface Counters (Fabric Interfaces) • Hot/Cold Interface Counters (Specific Interfaces)	• Interfaces • Services • Virtual Networks • Routing Zones
Fabric Health SLE	• BGP • BGP Monitoring • EVPN • EVPN Host Flapping • EVPN VXLAN Type-3 Route Validation • EVPN VXLAN Type-5 Route Validation • VXLAN Flood List Validation • ECMP • ECMP Imbalance (Fabric Interfaces)	• Services • Virtual Networks • Routing Zones
Virtual Infra Health SLE	• Hypervisor and Fabric VLAN Config Mismatch • VMs without Fabric Configured VLANs • Hypervisor Redundancy Checks	• Services

Table 87: Probes to be Enabled in Apstra (*Continued*)

SLE	Probe	Affected Items
Service Health SLE	• Device System Health Probe • Drain Traffic Anomaly Probe • Device Environmental Checks • Optical Transceivers • Interface Flapping (Fabric Interfaces) • Interface Flapping (Specific Interfaces) • Hot/Cold Interface Counters (Fabric Interfaces) • Hot/Cold Interface Counters (Specific Interfaces)	• Services • Virtual Networks • Routing Zones

RELATED DOCUMENTATION

[About the Service Levels Page](#) | 333

About the Service Levels Page

IN THIS SECTION

● [Tasks You Can Perform](#) | 334

To access this page, click **Monitor** > **Service Levels**.

You can view SLE information at the organizational level, site level, or at the device level. Select **Entire Org**, **Site** or **System** to view the respective SLE scores. Alternatively, you can use the search box to search for a specific site or a device to view more detailed information about the service levels at the site or device-level. You can select a time interval to get information about changes in SLE scores during a specific period.

- **SLE at the organization-level**—Provides a high-level view of the service levels for the entire organization. You can view information about all the sites in the organization or the worst performing sites in the organization and analyze what is causing deterioration to the service levels.
- **SLE at the site-level**—Displays SLE values for the selected site. The two tabs provide information about the system Health SLE and the Link Health SLE. It also provides information about the various classifiers and sub-classifiers that are contributing to the SLE score. Click the SLE or the classifiers to get more detailed insights in to the SLE score.
- **SLE at the system-level**—Provides information about the SLE is contributing to the SLE score at the device-level. Click the SLE or the classifiers to view information about the percentage of time a classifier contributed negatively or positively to the SLE score. For example, if bad optics caused degradation of performance, the value or percentage contributed by the Bad Optics classifier will be displayed on this page. As an administrator, you can use this information troubleshoot problems that are degrading the SLE score.

Tasks You Can Perform

- **View System Health SLE**—The System Health SLE continuously monitors the traffic on the device, any change in configuration, the hardware components, and CPU, memory and disk utilization by the device. A change in any of these parameters result in a change in the SLE score.
- **View Link Health SLE**—The Link Health SLE provides information about how the interface classifiers are contributing to the SLE score.
- **View Fabric Health SLE**—The Fabric Health SLE provides information about the health of physical and the virtual fabric infrastructure. A decrease in SLE score indicates issues with BGP sessions, route availability between leaf, spine, and superspine, route selection, and so on.
- **View Virtual Infra Health SLE**—The Virtual Infra Health SLE provides information about deviation from intended behavior in the virtual infrastructure. A decrease in the classifier value indicates issues with VLAN configuration or absence of redundant hypervisors.
- **View Service Health SLE**—The Service Health SLE monitors and provides information about the overall service performance impact due to system health and link health related anomalies.
- **View Overall Service at the organization Level**—You can get an overall view of the service levels at the organizational level. The tab displays the aggregated score for all the sites and for each of the SLEs.

- **Analyze SLE Score**—View more information about the circumstances that led to a lower SLE score. The Root Cause Analysis page displays additional information such as timeline, distribution, affected items and so on for the selected classifier. Click an SLE or its classifiers to view the root cause analysis information.

18

CHAPTER

Alerts

IN THIS CHAPTER

- [About the Alerts Page | 337](#)
 - [Configure an Alert Template | 338](#)
 - [Configure E-mail Notification for Alerts | 338](#)
-

About the Alerts Page

IN THIS SECTION

- [Tasks You Can Perform](#) | 337

To access this page, click **Monitor** > **Alerts**.

The Alerts page displays the alerts generated to notify administrators about anomalies in the network. To monitor specific alerts, you can apply an alert template for your organization. Alert templates filter the alert list to display only those alerts that are tracked in the template. You can also choose to receive e-mail notifications for the alerts.

The widgets display the total number of critical alerts that need immediate attention, warning alerts, and informational alerts.



NOTE: To receive e-mail notifications, you must configure Webhooks on the **Organization** > **Settings** page.

Tasks You Can Perform

- View alerts—View Alerts for the selected period and their severity for the entire organization or for specific sites in the organization.
- Configure e-mail notification for alerts—Configure to send email notification to site administrator or organization administrator when an alert is generated.
- Create an alert template—Create an alert configuration template that you can apply at the site or organization-level.
- Download alert details in a CSV file—Click the download button next to **Alerts Configuration**.

Configure an Alert Template

Juniper Data Center Assurance allows you to create alert templates to notify administrators about specific event types. You can apply an alert template to an organization to display only a filtered list of alerts on the Alerts page or send email notifications. If an alert template is not configured, all generated alerts are displayed on the Alerts page.

To create an alert template:

1. Log into Juniper Data Center Assurance.
2. Click **Monitor > Alerts > Alerts Configuration**.
3. Click **Create Template**.
 - a. Specify whether the template applies to the entire organization or for specific sites only.
 - b. Specify a template name.
 - c. Specify who should receive e-mail notifications (organization administrators, site administrators, or other specified e-mail addresses)
 - d. Select **Enable Alert** and select the alert types that should be displayed on the Alerts page, and **Email Notification** to send email notification for the selected alerts.
4. Click **Create**.

The alert template is created.

Configure E-mail Notification for Alerts

As an administrator, you can configure Juniper Data Center Assurance to send alert notification through e-mail. This enables administrators to respond to events quickly. You also create an alert template and apply at the site-level or at the organization-level.

Follow these steps to configure alert notification:

1. Log into Juniper Data Center Assurance.
2. Click **Monitor > Alerts**.

The Alerts page is displayed.

 - Select **To organization admins** to send e-mail notifications to administrators of the organization
 - Select **To site admins** to send e-mail notifications to administrators of the site
3. From the Alert Types, select the alerts that you want to be notified in e-mail.

- Enable Alerts—Select to display alerts on the Alerts page.
- Send Email Notification—Select to send e-mail notifications.

4. Click **Save**.

You will receive an e-mail notification when an alert is generated.

19

CHAPTER

Webhooks Integration

IN THIS CHAPTER

- [Configure Webhooks to Receive Event Notifications in Third-Party Applications | 341](#)
 - [Webhooks Source Address | 344](#)
 - [Webhook Delivery Events | 345](#)
 - [View Webhooks Delivery Events | 347](#)
-

Configure Webhooks to Receive Event Notifications in Third-Party Applications

You use webhooks to automate sending event notifications from a source application to a destination application. You can configure webhooks to enable Juniper Data Center Assurance to send notifications to third-party applications, such as ServiceNow, when events you have subscribed to are triggered on the managed devices.



IMPORTANT: Before you begin, ensure your firewall allows inbound traffic from the webhook source IPs.

Webhook source IPs are addresses used by the application to deliver webhook notifications to your configured endpoint. For information on webhook source IPs, see "[Webhooks Source Address](#)" on page 344.

To receive webhook notifications in the required format, you need to configure an intermediary that can interact with the sending and receiving applications. The recommended intermediary platform is Make. To process notifications, Make uses a workflow called Scenario, which converts the notifications to the desired format. Each event notification is sent to a URL that is generated for the Scenario in Make. The notification is then converted into a format that the application supports and is then delivered to the application.

For information on Scenarios in Make, see [Scenarios](#).

To configure webhooks to send notifications:

1. Log in to Make, <https://www.make.com/en/login>. From the home page, navigate to Scenario on the left navigation menu.

2. Configure the scenario settings as described, see [Create Your First Scenario](#).

Make generates a URL. Whenever an event is triggered, webhook notifications are sent to this URL.

3. Navigate to Organization Settings (**Organization > Settings**).

The Organization Settings page appears.

4. In the Webhooks tile, click **Add Webhooks**.



NOTE: You must be a superuser to configure webhooks.

The Add Webhook page appears.

5. Configure the webhooks settings. See [Table 88 on page 343](#) for webhooks field descriptions.



NOTE: In the URL field, enter the URL generated in step 2.

6. Click **Add** to configure the webhook.

You can find the webhook listed under the **Webhooks** tile.

7. Click **View** to see the delivery details of subscribed events.

Webhooks

Add Webhook

Name	⬆	Status	No. Topics	Deliveries
eventdops	⬆	✓	10	View

Use the Webhooks Deliveries page to verify the end-to-end connectivity between your organization and the configured webhook endpoints. For more information, see ["View Webhooks Delivery Events" on page 347](#).



NOTE:

- You must have access to the third-party application to view the event notifications.
- You must be an administrator with the Network Admin role to perform corrective action for the notification received.

Table 88: Parameters to Configure Webhooks

Field	Description
Status	Select to enable or disable webhooks. The values are: • Enabled: Webhooks is enabled and you can get notifications on third-party applications when events you have subscribed to occur in the organization. • Disabled: Webhooks is not enabled and you cannot get notifications on third-party applications when events occur in your organization.
Webhook Type	Select the format in which notifications are to be sent when a subscribed event occurs. The options are: • HTTP POST • OAuth 2.0 • Splunk
Name	Enter the name of the server or application to which notifications for subscribed events are to be sent.
URL	Enter the URL of the server or application where the notifications are to be sent when a subscribed event occurs. NOTE: Ensure the URL endpoint is publicly accessible. You must configure webhooks to send notifications to third-party applications, such as ServiceNow, when events you have subscribed to are triggered on the managed devices. To receive webhook notifications in a compatible format, you need to configure an intermediary that can interact with the sending and receiving applications. The recommended intermediary platform is Make.
Topics	Select the types of events for which you want to receive webhook notifications. Select Audits or Data Center Events to receive notification when an audit log or data center event is generated.
Advanced Settings	
Verify Certificate	Enable or disable verification of certificates.

Table 88: Parameters to Configure Webhooks *(Continued)*

Field	Description
Secret	Enter the secret to validate that the notifications received are from valid hosts.
Custom Headers	
Key	Enter a unique key that the webhook endpoint can use to authenticate the event notifications.
Value	Enter a unique value for the key.

Webhooks Source Address

SUMMARY

This topic lists the source IP subnets that must be allowed on your firewall to receive event notifications through webhooks.

IN THIS SECTION

- [Webhook Source IP Addresses for Juniper Cloud Instances | 344](#)

Webhook Source IP Addresses for Juniper Cloud Instances

To successfully receive event notifications through webhooks, you must configure your firewall to allow incoming traffic from the webhook source IP subnets for your cloud instance. Webhook source subnets are IP address ranges used by the application to deliver webhook notifications to configured endpoint.

Webhook event delivery may fail if your firewall blocks traffic from any IP addresses within these subnets.

When an event occurs, a webhook notification is sent to the configured endpoint. To ensure notifications are delivered to the configured endpoint, configure your firewall to allow inbound traffic from the following subnets.

- 34.0.0.0/8
- 35.0.0.0/8
- 44.0.0.0/8

- 52.0.0.0/8
- 54.0.0.0/8
- 23.0.0.0/8



TIP: If you are not receiving webhook events:

- Ensure that your firewall is not blocking traffic from any IP addresses within the subnets listed above.
- Verify that the webhook endpoint URL you configured is publicly accessible and can accept incoming HTTP/HTTPS traffic.

Webhook Delivery Events

You can verify the end-to-end connectivity between your organization and a webhook endpoint from the Webhook Deliveries page. When an event occurs, a webhook notification is sent to the configured webhook endpoint.

Use the Webhook Deliveries page to monitor and review the status of the attempted webhook deliveries. See, "[View Webhooks Delivery Events](#)" on page 347.



REMEMBER: To view webhook delivery events, you must first configure webhooks and select the topics (events) for which you want to receive notifications. The webhook topics that support webhook delivery are Alerts, Audits, Pings, and Device Up/Down.

For more information, see "[Configure Webhooks to Receive Event Notifications in Third-Party Applications](#)" on page 341.

If you want to test a webhook but do not have a webhook endpoint available, you can use public webhook receivers that enable you to receive and inspect the incoming webhook messages.

These public sites generate a temporary URL where you can send webhook requests.

Examples include:

- <https://webhook.site/>
- <https://public.requestbin.com/>



CAUTION: Note that these are third-party webhook receiver sites. Exercise caution when sharing sensitive information with these sites.

Make sure to delete any test webhooks from the organization if they are no longer in use. Otherwise, the organization instance will continue to push webhook notifications to the endpoint and the the webhook receiver might blacklist the IP address of the instance.

Once a webhook is configured, you can use the Webhooks Deliveries page to monitor and troubleshoot webhook deliveries. From this page, you can view detailed information about all attempts made by the system to send notification to the configured endpoint. The information includes timestamps, headers, request and response body, HTTP status codes, error messages, and so on.

You can use the Webhooks Deliveries page to:

- Debug failed webhook attempts
- Inspect payload content
- Monitor webhook performance over time
- Identify authorization failures and rate-limiting issues
- Search webhook events by topic, error message, or status code
- Filter webhook events by time interval and delivery status
 - To filter based on time interval you select, click *Time Interval* (by default it is set to Today). You can choose Last 60 Minutes, Last 24 Hours, Last 7 Days, Custom Date (enter a custom date) Today, Yesterday, This Week, or Custom (enter a custom time range).
 - To filter based on webhook delivery status, click a status filter.

Click **Total** to view all events, **Good** to view successfully delivered events, and **Bad** to view failed to deliver events.

The rows are color-coded to provide a quick visual indicator of the webhook status. A green bar indicates a successful delivery and a red bar indicates a failed delivery.

Delivery Events			
3 Total 1 Good 2 Bad			
Date	Topic	Status	Error Message
11:46:30.735 AM Jan 15, 2025	Ping	Success (200)	
11:46:30.829 AM Jan 15, 2025	Ping	Failure (N/A)	throttled due to intern.
11:46:30.370 AM Jan 15, 2025	Ping	Failure (502)	URL: https://api.webhc Error: 502 Bad Gatewa

View Webhooks Delivery Events

Use the Webhook Deliveries page to monitor and review the status of the attempted webhook deliveries.

To view webhook delivery events:

1. Navigate to **Organization > Settings** and locate the Webhooks tile.
2. Click **View** next to the desired webhook.

Webhooks				
				Add Webhook
Name	⬆	Status	No. Topics	Deliveries
eventdops	⬆	✓	10	View

The Webhook Deliveries page appears.

3. View the webhook event notifications in the Webhook Deliveries overlay page.

For field description, see [Table 89 on page 348](#).



NOTE: Webhook delivery events are retained for up to 61 days.

You can click an event to view detailed information about the event.

Webhook Deliveries

Delivery status of recent webhooks for eventdops

Today

Delivery Events

458 Total 0 Good 458 Bad

Status	Error Message
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests
Failure (429)	URL: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc Error: 429 Too Many Requests

Timestamp

6:09:56.783 PM Jul 14, 2025

Topic

Audits

Status Code

429

Error Message

url: https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc failed: 429 Too Many Requests

Request URL

https://webhook.site/cbf953d8-c7f3-4645-8aec-a77b2b215fdc

Request Headers

Content Type: application/json
User Agent: Mist-webhook

Request Payload

```
{
  "topic": "audits",
  "events": [
    {
      "admin_name": "jcloud monitoring jcloud-
monitoring@juniper.net",
      "id": "7587cae7-ea7e-4168-915c-98df31226578",
      "message": "Accessed Org \"Main Stage Org\"",
      "org_id": " ",
      "src_ip": " ",
      "timestamp": "1752496795.429254",
      "transparent": "00-0e16a1d7728d9db62b3e5ab74b123c6f-
c72f879b36393152-01",
      "user_agent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/125.0.0.0
Safari/537.36"
    }
  ]
}
```

Response Headers

Content Type: text/html

For more information on the fields, see [Table 90 on page 349](#).

4. Click **X** to exit the overlay.

Table 89: Webhook Event Field Description

Field	Description
Date	Date and time when the webhook delivery for the event was attempted.

Table 89: Webhook Event Field Description (*Continued*)

Field	Description
Topic	Type of event that triggered the webhook. The webhook topics that support webhook delivery are alerts, audits, pings, device status. NOTE: When you create webhooks, you must select the topics that you want to be notified about.
Status	HTTP status code that indicates the result of the webhook delivery. For example: 200 (Success), 400 (Bad Request), 429 (Too Many Requests).
Error message	Provides reason for the failed webhook delivery, if applicable.

Table 90: Webhook Details Pane

Field	Description
Timestamp	Date and time when the webhook delivery for the event was attempted.
Topic	Type of event that triggered the webhook. The webhook topics that support webhook delivery are alerts, audits, pings, device status. NOTE: When you create webhooks, you must select the topics that you want to be notified about.
Status Code	HTTP status code that indicates the result of the webhook delivery. For example: 200 (Success), 400 (Bad Request), 429 (Too Many Requests).
Error Message	Provides reason for the failed webhook delivery, if applicable.
Request URL	Address of the configured endpoint where the webhook was sent.
Request Headers	Metadata included with the webhook request.

Table 90: Webhook Details Pane *(Continued)*

Field	Description
Request Payload	Displays the requested information in JSON format. The request body includes detailed description of the event.
Response Headers	Metadata included with the response.

20

CHAPTER

Help & Support

IN THIS CHAPTER

- [Provide Feedback About DC Assurance | 352](#)
 - [About the Support Tickets Page | 354](#)
 - [Create a Support Ticket | 360](#)
 - [Requesting Technical Support | 370](#)
-

Provide Feedback About DC Assurance

You can provide feedback about your experience with Juniper Data Center Assurance using the feedback form available on the GUI. Your feedback, comments, and suggestions help us improve your experience and understand your expectations from the product.

To provide feedback about Juniper Data Center Assurance:

1. Click the **Feedback** icon on the top right corner of the GUI.

Figure 108: Feedback Icon



The **Help Us Improve** dialog box opens.

2. Enter your feedback in the dialog box.

Figure 109: Help Us Improve

The screenshot shows a feedback form titled "Help Us Improve" with a close button (X) in the top right corner. The form is for "Apstra Cloud Services feedback".

The first section asks, "Rate your first impression of the product? *". It features five icons representing a rating scale: a thumbs down (Worst), a sad face (Not Good), a neutral face (Fine), a happy face (Looks Good), and a thumbs up (Very Good). A horizontal slider bar is positioned below these icons.

The second section asks, "What was the most useful feature to you? *". It includes five radio button options: Dashboard, Marvis, Service Awareness, Impact Analysis, and Service Level Expectations.

The third section asks, "What would you like to explore next?". It contains two buttons: "New Feature Request" (highlighted in blue) and "Issue" (grey). Below these buttons is a large text input field with the placeholder text "Enter your New Feature Request...".

The fourth section is titled "Apstra feedback" and has a toggle switch that is currently turned on. It also asks, "What would you like to explore next?". It features the same "New Feature Request" and "Issue" buttons, followed by another large text input field with the placeholder text "Enter your New Feature Request...".


At the bottom right of the form, there are two buttons: "Submit" (blue) and "Cancel" (grey).

- Provide your rating for the application by using the icons or the slider.
- Select the feature that you find the most useful.
- To request for a new feature on DC Assurance, click the **New Feature Request** button and enter the description in the text field.
- If you are facing any issues on DC Assurance, click the **Issue** button and provide details in the text field.

- e. Enable the toggle for **Apstra feedback** to submit a feature request or to raise an issue regarding Apstra Data Center Director.
3. Click **Submit** after entering information in the mandatory fields.

The following message is displayed:

Figure 110: Feedback Submitted Message




About the Support Tickets Page

IN THIS SECTION

- [Tasks You Can Perform | 355](#)
- [Field Descriptions | 357](#)

The Support Tickets page enables you to easily report issues, get assistance, and track the status of the issues directly within the application. The page enhances your overall support experience by providing seamless access to Juniper support services, whenever you need them.

To access the Support Tickets Page, navigate to the **Help menu** (?) on the top-right corner of the banner and click **Support Tickets**. You can view, create, and track support tickets associated with your organization from this page.




Tasks You Can Perform

- View details about all the support tickets that are associated with your organization, such as ticket number, ticket type, ticket summary, ticket status, and so on.

Click a support ticket to view detailed description of the support ticket. For more information, see [Table 92 on page 358](#).

- Create a support ticket. For more information, see ["Create a Support Ticket" on page 360](#).
- Update a support ticket. For more information, see ["Add a Comment" on page 359](#).
- Resolve a support ticket.

To resolve a support ticket, select an Open or Pending ticket and click **Mark Ticket as Solved** in the Support Tickets Details page.



A confirmation pop-up appears. Click **Mark as Solved** to close the ticket.

The ticket status changes to Solved.

- Navigate to the Juniper support page to find relevant resources and documentation that can help you resolve the issue on your own.

Click the link provided on the page to be directed to the Juniper support page.


- Filter support tickets displayed in the table by time interval.

To filter tickets based on time interval you select, click *Time Interval* (by default, it is set to This Month). You can choose Last 24 Hours, Last 7 Days, Last 30 Days, Last 12 Months, Custom Date (select a custom date) Today, This Week, This Month, This Year, or Custom Range (enter a custom time and date range).



NOTE: The notification badge on the **Help menu** (?) in the top-right corner displays the number of unread support tickets for the logged-in user. The ticket count gets updated when you open an unread ticket. It also gets updated when you select a different time-range. For example, if you select **Last 7 Days**, the ticket count gets updated to show the

number of your unread tickets from the last seven days.



- Click **Refresh** icon to manually refresh the Support Tickets page.

Field Descriptions

Table 91 on page 357 displays fields on the Support Tickets page.

Table 91: Fields on the Support Tickets Page

Field	Description
Ticket Number	Serial number associated with the support ticket.
Ticket Type	Type of ticket. The options are: • Questions • Configuration Help • Problem
Ticket Summary	Short description of the issue.

Table 91: Fields on the Support Tickets Page (Continued)

Field	Description
Created	Date and time when the support ticket was created.
Last Update	Date and time when the support ticket was last updated.
Requested by	Name of the user who created the support ticket.
Status	Status of the support ticket. The options are: • Open—The ticket is created and is in Open state. • Pending—The ticket resolution is in progress and is awaiting a response or action from the user. • Solved—The issue is resolved. If you are satisfied with the resolution, click Mark Ticket as Solved to resolve the ticket. • Closed—The ticket is closed. The status changes to Closed when the JTAC team closes the ticket.

Table 92: Fields on the Support Tickets Details Page

Field	Description
Ticket Number	Serial number associated with the support ticket.

Table 92: Fields on the Support Tickets Details Page *(Continued)*


Field	Description
Ticket Status	Status of the ticket. The options are: • Open—The ticket is in open state. The ticket remains open until the issue is resolved. • Solved—The issue is resolved. Click Mark Ticket as Solved to close the ticket. • Closed—The ticket is closed.
Information	
Requested By	Name of the person who created the support ticket.
Created	Date and time when the support ticket was created.
Last Update	Date and time when the support ticket was last updated.
CC (Optional)	Enter the e-mail IDs of people you wish to notify about the ticket. You can update this field even after ticket creation. To update the CC list, click an Open or Solved ticket and enter the e-mail IDs in the CC field, then click Save. This field is disabled when the ticket is Closed.
Description	Detailed summary of the issue, such as information about ticket type, impacted sites and devices, and other relevant information provided at the time of creating the ticket.
Comments	Add and view comments on the progress of the support ticket. To add a comment, click Add a Comment, enter your comment and click Submit. This section also displays comments from the JTAC team about the progress of the support ticket.

Create a Support Ticket


You can create a support ticket to get assistance from the Juniper Networks Technical Assistance Center (JTAC) team to resolve problems with your Juniper Networks products and services. Users with superuser or network admin roles can create a support ticket.

To create a support ticket:

1. Navigate to the **Help menu** (?) on the top-right corner of the banner and click **Support Tickets**.



NOTE: The notification badge on the **Help menu** (?) in the top-right corner displays the number of unread support tickets for the logged-in user. The ticket count gets updated when you open an unread ticket. It also gets updated when you select a different time-range to filter the support tickets. For example, if you select **Last 7 Days**, the ticket count gets updated to show the number of your unread tickets from the last seven days.



The Support Tickets page appears.

2. Click **Create a Ticket**.

The Support Tickets: New Ticket page appears.


3. Select the **Technology** for which you are creating the ticket.

The options are:

- Routing Assurance
- Apstra
- JSI

4. Select the **Ticket Type**.

The ticket types are:




A screenshot of a web form titled "Ticket Type". Below the title is a dropdown menu currently showing "Questions". The dropdown is open, displaying three options: "Questions", "Configuration Help", and "Problem".

- **Questions**—Select this ticket type if you have a general question.
- **Configuration Help**—Select this ticket type if you need help with configuring your organization, site, network, or devices.
- **Problem**—Select this ticket type if you need help with issues like an outage, setting up a new site, or returning a merchandise. You must also select the impact from the Impact drop-down list.

The Impact drop-down list allows you to select a more specific problem type or identify the entities impacted by the problem.

Impact options are:



A screenshot of a web form showing the "Ticket Type" dropdown set to "Problem". Below it, the "Ticket Summary" field is visible. To the right, the "Impact" dropdown menu is open, showing a list of options: "Specific Devices Impacted", "Partial Network Impacted", "New Site Bring Up", "RMA", and "Full Production Network Impacted". The "Specific Devices Impacted" option is highlighted with a black border.

- **Specific Devices Impacted**—The problem is limited to specific devices on the network.
- **Partial Network Impacted**—The problem affects part of your network but not your entire network.
- **New Site Bring Up**—Problems in bringing up a new site that is not in production yet.
- **RMA**—You want to request a Return Material Authorization to repair or replace a defective device.

- **Full Production Network Impacted**—The problem affects your entire network.
5. Before you proceed to create a support ticket, you can try to resolve the issue with the help of our AI-powered conversational assistant, Marvis. Start a conversation with Marvis and get instant assistance with resolving your issues.

Marvis scans a range of documentation and generates tailored responses to help you resolve issues. You can ask follow-up questions and request clarifications to refine the search results. Additionally, Marvis suggests related questions and provides links to Juniper resources based on which the response is generated.



Video: [AI-assisted case resolution with Marvis](#)

The **How can I help you today?** section offers a set of prompt templates to help you start your conversation with Marvis. Alternatively, you can type your questions directly into the **Ask Marvis a question** text box and Marvis will assist you.

To use a prompt template, click the one that best suits your question.

- **I need help troubleshooting...**—Use this when you are trying to diagnose or resolve an issue.
- **How do I configure...**—Use this when you need guidance on setting up or modifying a configuration.
- **I have an issue with...**—Use this to describe a specific problem you are encountering.
- **Can you explain how...**—Use this if you want to learn or need clarification about a feature or functionality.




NOTE: These prompts are just suggestions and are not meant to limit what you can ask Marvis. As long as the question is about Juniper products and services, there are no constraints to the range of questions you can ask Marvis.

A textbox appears on the right side of the page with the selected prompt auto-populated. Complete the prompt with the relevant details and press **Enter**. Based on your input, Marvis will instantly provide relevant responses.

If your question doesn't fit the suggested prompt templates, you can ignore the suggestions and enter your question directly in the **Ask Marvis a question** textbox. Press **Enter** to generate a response. Marvis will use your input to search for relevant solutions and assist you instantly.

An action toolbar appears at the end of each response, providing quick access to the following functions:



- **Copy to Clipboard**—Click the copy icon to copy your question, the generated response, and related citations to the clipboard.
- **Provide Feedback**—Share your experience resolving the issue with the help of Marvis.

Based on your experience, click the Thumbs Up or Thumbs Down icon to rate your interaction with Marvis. The Marvis Feedback dialog box appears.

You are provided with a set of predefined feedback options, select the ones that best describe your experience. You can also add comments to offer suggestions or provide additional context for your feedback.

Figure 111: Feedback Dialog Box for Positive Feedback

Marvis Feedback

What did Marvis do well?

Accurate and helpful

✓

Clear and concise

✓

Well-structured response

Provided relevant information

Understood context well

✓

Offered insightful analysis

✓

Comment (max 1000 characters)

Submit

Figure 112: Feedback Dialog Box for Negative Feedback

Marvis Feedback

How can Marvis improve?

Inaccurate information	Unclear or confusing
Irrelevant to the question	Incomplete response ✓
Misunderstood context ✓	Poorly structured

Comment (max 1000 characters)

Submit

Click **Submit** to submit your feedback.



NOTE: Before you submit, review your feedback carefully as you cannot reset or modify your feedback once submitted.

- **View Citations**—Click the citations icon to view the list of sources on which the generated response is based. Click the citation links to explore the source materials.
- **Reset Conversation**—Click the reset conversation icon to start a new conversation.



NOTE: Resetting the conversation clears your conversation history and starts a new session.

Below the action toolbar, you can find a list of related questions, that can help you refine your query and explore additional information. Click a question from the list to populate it in the chat and continue the conversation or type your question in the textbox.

If you are able to resolve the issue with the help of Marvis, you can cancel the support ticket by clicking **Cancel** at the top-right corner of the New Ticket page.

6. If your issue needs deeper investigation and require assistance from the support team to resolve the issue, click **I still need to create a ticket** to create a support ticket.



NOTE: Mandatory fields are highlighted in red and may vary based on the technology you selected.

◀ Support Tickets : **New Ticket**

Description is required

Email attachments to support@juniper.net

Technology

☐ Routing Assurance
 ☐ Apstra
 ☒ JSI

Ticket Type

Questions ▾

Ticket Summary

How to configure BGP?

CC (optional)

Email addresses

Impacted Sites Add Site

Primary Site (07612eb5-0c50-43d1-bbce-aa3dc3b4560f)

Impacted Devices Add Device

Description

Submit Ticket

7. Provide a short description of the issue, details about the impacted sites, devices, and so on.
For field descriptions, see [Table 93 on page 368](#).



NOTE: You can send relevant logs and screenshots related to the ticket to support@juniper.net.

8. Click **Submit Ticket**.

On successful submission, the new ticket is listed on the Support Tickets page. You also receive a confirmation e-mail with the ticket details.

You can track the progress of the support ticket from the Support Tickets page. You and anyone copied on the ticket will also receive e-mail notifications whenever the ticket is updated.

Table 93: Ticket Summary Field Descriptions

Field	Ticket Types	Description
Ticket Summary	All Ticket Types	Enter a short description of the issue.
CC (Optional)	All Ticket Types	Enter the e-mail IDs of people you wish to notify about ticket updates.
Impacted Sites	All Ticket Types	Add the sites impacted by the issue. Click Add Site to add one or more sites or site groups impacted by the issue. Alternatively, you can manually enter the site details.
Impacted Devices	All Ticket Types	Select the devices impacted by the issue from the list of devices associated with the site or organization. NOTE: For Routing Assurance and JSI, this field is mandatory if there are onboarded devices in your organization. This field is optional if your organization does not have onboarded devices and you require support for non-onboarded devices. Click Add Device to add the impacted device. You can also search for the devices by entering the device name or MAC address.
Registered Edges	All Ticket Types	Select the edge device impacted by the issue. Click Add Registered Edge to add the impacted edge device. You can also search for the device by entering the device name.

Table 93: Ticket Summary Field Descriptions (*Continued*)

Field	Ticket Types	Description
Description	All Ticket Types	Enter a detailed description of the issue and any contextual information that can help the support team troubleshoot the issue. NOTE: • This field is mandatory if you select the Ticket Type as Questions. • Provide all the relevant details to ensure faster resolution. • You can send relevant logs and screenshots related to the ticket to support@juniper.net.
Time of Issue	• Configuration Help • Problem	Select the date and time when the issue occurred. To select the data and time, click the Time of Issue field and use the calendar and time list to select the date and time when the issue occurred. NOTE: By default, it is set to the current date and time.
Contact Number	Problem	Enter the contact number to receive calls from the JTAC team regarding the support ticket.
Juniper Account Team Name and Contact	Problem	Enter the name and contact details of your Juniper account representative. NOTE: This field is available only if you select Full Production Network Impacted from the Impact drop-down list.

Requesting Technical Support

IN THIS SECTION

- Self-Help Online Tools and Resources | 370
- Creating a Service Request with JTAC | 371

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the [JTAC User Guide](#).
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal, Customer Support Center (CSC), that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions to your questions using our Knowledge Base: <https://kb.juniper.net/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://support.juniper.net/support/requesting-support/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool:
<https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://support.juniper.net/support/requesting-support/>
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico)

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/>.