

Juniper Mist Network Monitoring and Troubleshooting

Published
2024-01-02

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 Mist Network Monitoring and Troubleshooting
Copyright © 2024 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

1

Overview

Introduction | 2

Requirements | 3

AIOps in Action | 4

Explore More Features | 12

2

Insights

Mist Insights | 14

3

Service Level Expectations

Service-Level Expectations Overview | 24

Wireless SLEs | 32

Time to Connect SLE | 33

Wireless Successful Connects SLE | 35

Coverage SLE | 37

Roaming SLE | 39

Wireless Throughput SLE | 41

Capacity SLE | 43

AP Health SLE | 44

Wired SLEs | 46

Wired Throughput SLE | 47

Switch Health SLE | 50

Wired Successful Connect SLE | 52

WAN SLEs | 54

WAN Edge Health SLE | 55

| WAN Link Health SLE | 57

Location SLEs | 59

| SDK Connect Time SLE | 60

| Location Latency SLE | 61

| Teleports SLE | 63

| Dropped Requests SLE | 64

| Location AP Health SLE | 67

4

Alerts

Alert Configuration | 70

| Alerts Overview | 70

Juniper Mist Alert Types | 74

1

CHAPTER

Overview

[Introduction](#) | 2

[Requirements](#) | 3

[AIOps in Action](#) | 4

[Explore More Features](#) | 12

Introduction

If your job involves troubleshooting problems, investigating user complaints, or tracking network performance, you'll find that all these tasks become easier with the AI-driven operations (AIOps) features in your Juniper Mist portal.

AIOps is embedded into Juniper Mist, enabling your IT operations team to stay on top of and manage all the complexity of your distributed networks. Mist AI applies big data, analytics, and machine learning capabilities to intelligently sift through network information to pinpoint events and recognize patterns that indicate potential issues. Mist AI can also diagnose the root cause of an issue and recommend action.

These features shorten the time spent on troubleshooting and empower you to take proactive actions to ensure positive user experiences. No more guessing about the scope of an incident. No more needle-in-a-haystack searches through log files to identify root causes. No more struggling to reproduce issues so that you can capture packets.

With the Juniper Mist dashboards, you'll see:

- Success/failure indicators that you can interpret at a glance
- Visualizations that show exactly when and where an issue originated
- Packet captures for every incident
- Root-cause analysis

And even better, you can discover many issues before they have an impact. With the Service Level dashboards, you can quickly spot any conditions that don't meet your expectations. Take action before incidents occur.

And if you have a Marvis Virtual Network Assistant subscription, you also get:

- AI-recommended actions to improve network performance and user experiences
- Conversational support with issue identification and troubleshooting
- Robust query language for more structured inquiries
- Proactive identification of potential issues

Ready to get started?

- Understand the requirements, including user permissions and subscriptions. See ["Requirements" on page 3](#).

- Follow a day in the life of an operations engineer to how you can put these features to work for you and your users. See ["AIOps in Action" on page 4](#).

Requirements

IN THIS SECTION

- [User Role | 3](#)
- [Subscriptions | 3](#)

Your access depends on your role in the Juniper Mist™ portal and the subscriptions that you've activated for your organization.

User Role

The following user roles can access monitoring information in the Juniper Mist portal:

- Super User
- Network Admin
- Observer
- Helpdesk
- Super Observer

Subscriptions

Your subscriptions determine the features that are available to you in the Juniper Mist portal.

- Base Subscription—With the base subscription, you can:

- View AI-driven insights and easy-to-interpret graphs for site events, client events, AP events, and more.
- Configure alerts to get notified when events happen in your Juniper Mist organization.
- With a subscription for location services or wireless, wired, or WAN networking, you can monitor service levels and investigate issues impacting user experiences.
- Marvis Virtual Network Assistant Subscription—With a [Marvis Virtual Network Assistant](#) subscription, you can:
 - Chat with your conversational network assistant to ask questions and troubleshoot issues.
 - Submit structured queries using Marvis Query Language.
 - View the Marvis Actions page, which identifies issues, presents a root cause analysis, and recommends actions.
 - Use the Marvis Windows and Android client.
 - Integrate Juniper Mist with apps such as Microsoft Teams, ChatGPT, Zoom, and more.

AI Ops in Action

IN THIS SECTION

- [Starting the Day with the Marvis Actions Dashboard | 5](#)
- [Troubleshooting Low Service Levels | 6](#)
- [Getting Help from the Marvis Conversational Interface | 9](#)
- [More Video Demos | 11](#)

Let's see how Oscar, an operations lead, uses the Juniper Mist portal to anticipate and respond to issues during a typical day.

NOTE: As you read about Oscar's experiences, you'll get a high-level introduction to many features in the Juniper Mist portal. You'll get more in-depth information later in this guide.


Starting the Day with the Marvis Actions Dashboard

Oscar always starts his day by looking at the Marvis Actions dashboard. On this dashboard, Marvis identifies actions that can improve the user experience. By following through on these recommendations, Oscar can address issues *before* users report an impact.

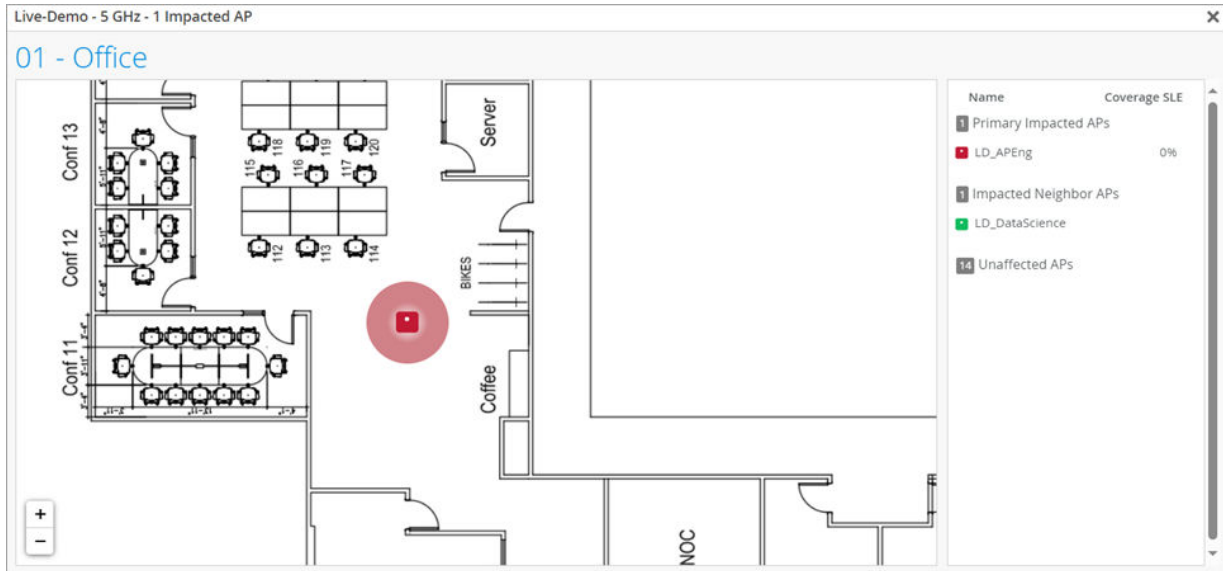
Today Oscar notices eight issues with APs. With one click, he sees a high-level root cause analysis: five are offline, one failed its health check, and one has a coverage hole.



He clicks the Coverage Hole item. At the bottom of the page, Marvis shows him where and when the issue occurred. Marvis also provides a recommendation to resolve the issue.

COVERAGE HOLE					
 RECOMMENDED ACTION The following APs noticed frequent coverage issues around them. Please reposition or add more APs in order to provide adequate coverage.					
<input type="checkbox"/>	Site	APs	Details	Date	Status
<input type="checkbox"/>	Live-Demo	LD_APEng	5 GHz View More	Nov 13, 2023 3:58 PM	Open

Oscar clicks to view more information. For this type of issue, Marvis displays the floorplan. Oscar sees exactly where this AP is located. With this information, he understands the issue and the impact and can follow through to ensure adequate coverage.



Video Demo

In this video demo, Marvis recommends actions for bad signal strength.

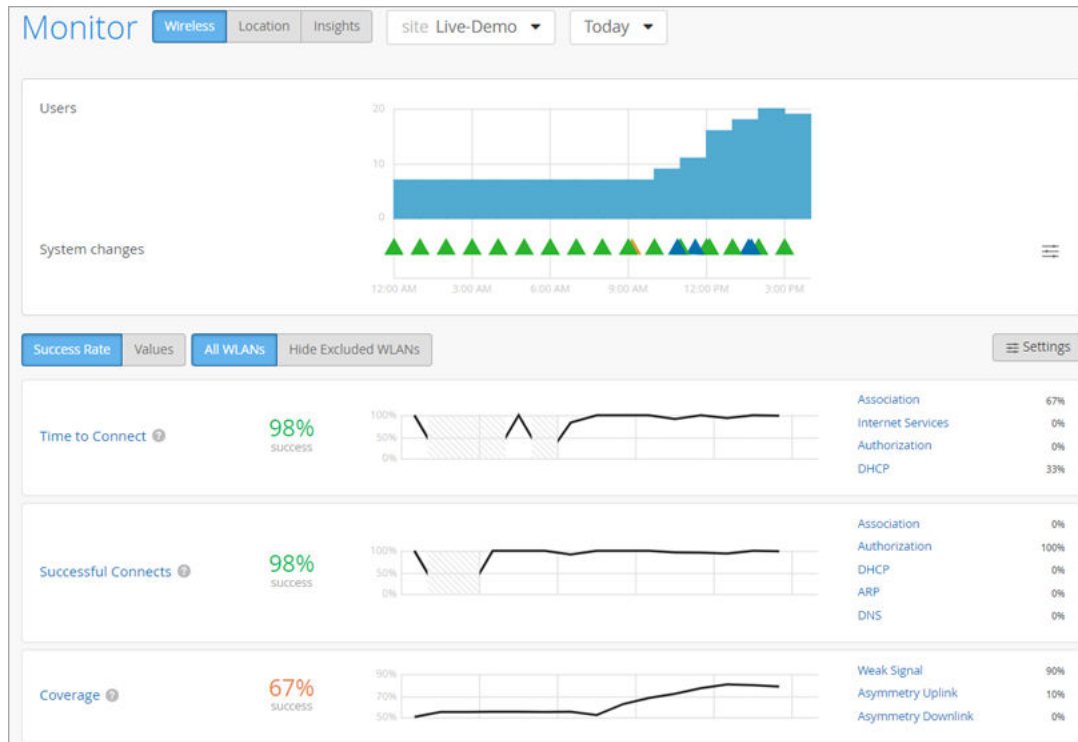


Video:

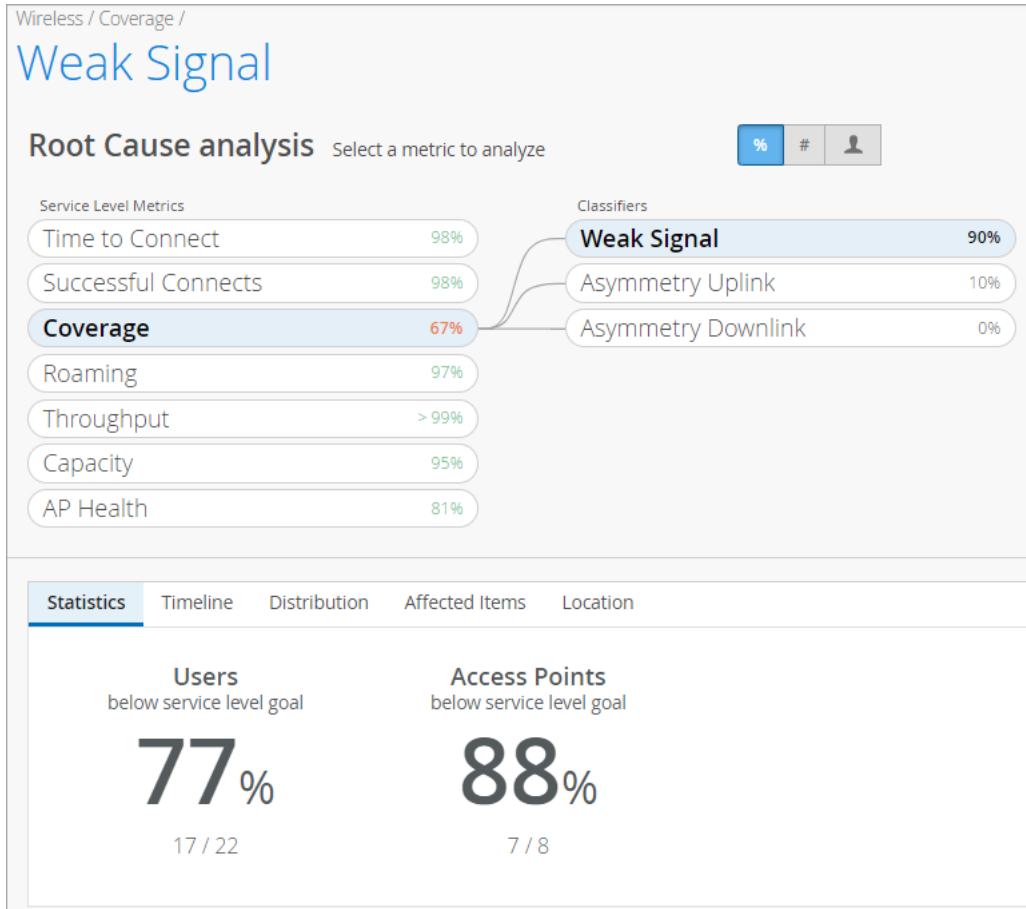
Troubleshooting Low Service Levels

Next, Oscar turns to the Service Level dashboards. These dashboards show successes and failures for critical factors (SLEs) that can impact user experiences.

On the Wireless dashboard, color coding draws Oscar's attention to a low SLE for coverage. On the left side, he sees the overall success rate for each service level. Coverage has only a 67 percent success rate. On the right side of the page, Oscar sees a high-level root cause analysis (on the right). Of the unsuccessful user experiences, 90 percent are due to weak signal.



Oscar clicks to take a closer look. On the Root Cause Analysis page, he clicks **Weak Signal** to view more information. He can see that 77 percent of users and 88 percent of APs are having signal issues.



By using the tabs in the lower half of the screen, Oscar can get a complete view of the scope of impact:

- Timeline—When did the issues occur?
- Distribution—Where in the network did the issues occur?
- Affected Items—Which users, devices, and applications were involved?
- Location—Where are the floorplan did the issues occur?

Video Demo

This video demo shows how to troubleshoot low SLEs for WAN issues.

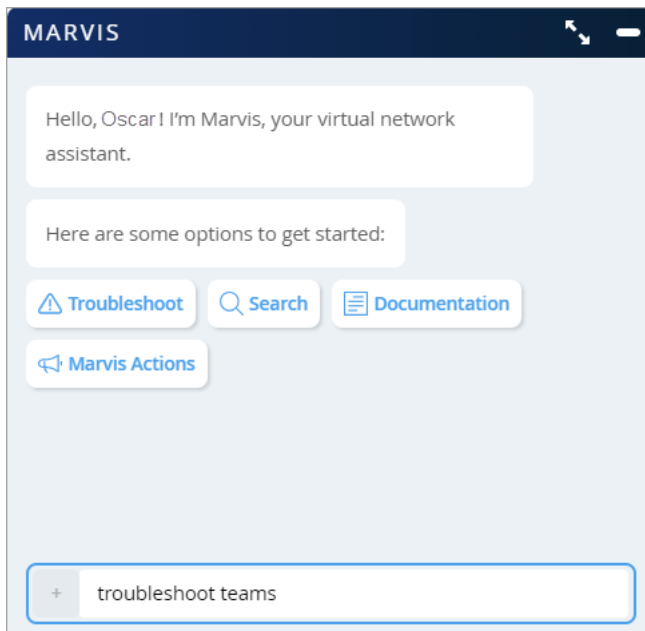


Video:

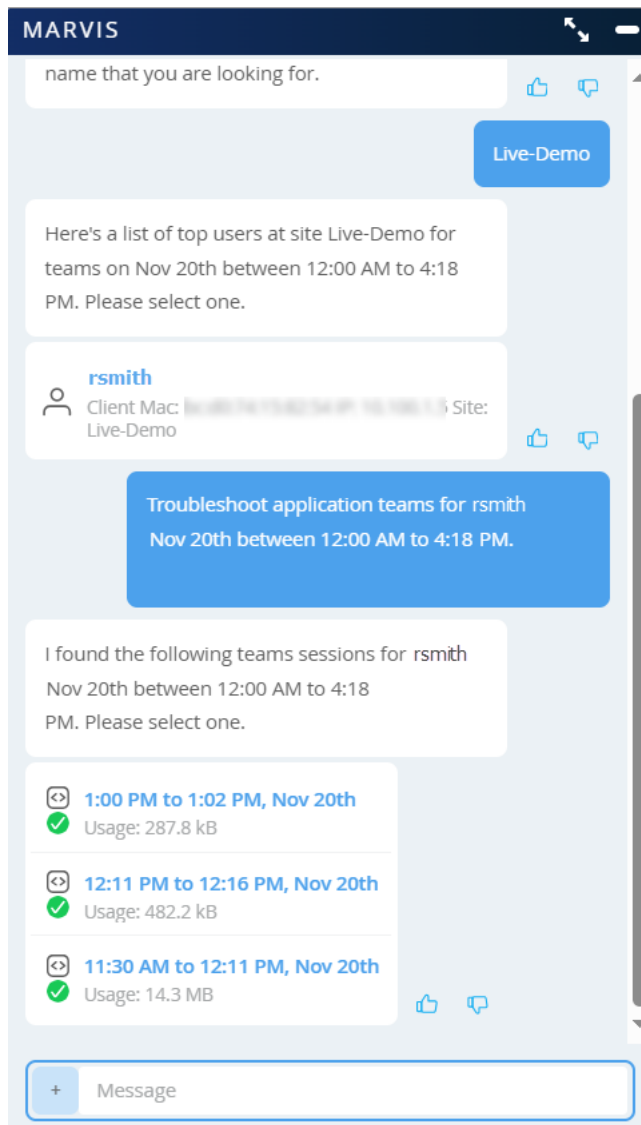
Getting Help from the Marvis Conversational Interface

After lunch, Oscar bumps into a colleague, Roberta, who mentions that she had a bad Microsoft Teams call that morning. Oscar decides to get help from Marvis by using the chat feature. He clicks the Marvis icon at the bottom left corner of the screen.

In the pop-up window, he enters: *troubleshoot teams*.

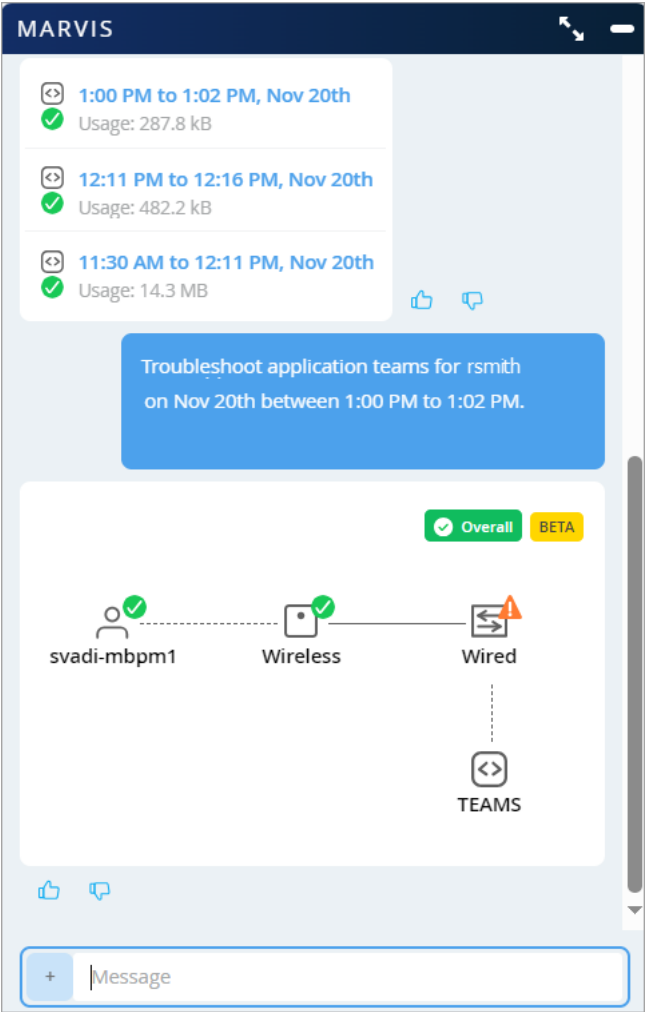


As Marvis asks guiding questions and Oscar replies, Marvis provides a list of recent Teams calls.



Oscar clicks a call to view additional information.

Marvis shows that there was an issue on the Wired network. From here, Oscar can click to view additional information.



Video Demo

In this video demo, Marvis helps to troubleshoot an issue with Microsoft Teams.



Video:

More Video Demos

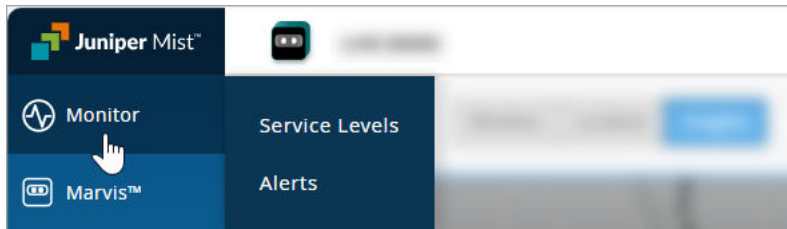
For even more insights into these features, explore more videos:

- [10-Minute Troubleshooting Video](#)
- [Everyone Does Client Health \(SLE v2\) Video](#)

Explore More Features

With this introduction, you're ready to start exploring the Juniper Mist™ portal on your own.

This guide covers the features that you can access through the **Monitor** menu.



- To get started with Service Levels, see:
 - ["Mist Insights" on page 14](#)
 - ["Service-Level Expectations Overview" on page 24](#)
- To get started with Alerts, see ["Alert Configuration" on page 70](#).

If you're a Marvis subscriber, you're likely to go back and forth between the **Monitor** menu and the **Marvis** menu to get all the benefits of AIOps features, insights, and assistance. To keep learning about Marvis, see the [Juniper Mist Marvis Guide](#).

2

CHAPTER

Insights

Mist Insights | 14

Mist Insights

SUMMARY

Get a high-level view of the the events that have occurred at each site in your organization. You can adjust the view to focus on any site or device.

IN THIS SECTION

- [What Data Are Used for Insights? | 14](#)
- [Finding the Insights Dashboard | 15](#)
- [Selecting the Context and Time Period | 15](#)
- [Layout of the Monitor Page | 16](#)
- [Map Display | 16](#)
- [Insights Timeline \(Time Range\) | 17](#)
- [Site Events | 17](#)
- [Client Events | 17](#)
- [AP Events | 20](#)
- [Applications | 20](#)
- [Network Servers | 20](#)
- [Pre-Connection and Post-Connection | 21](#)
- [Current Site Properties | 21](#)
- [Current WLANs | 22](#)
- [Access Points | 22](#)
- [Clients | 22](#)
- [Wired Switches | 22](#)

What Data Are Used for Insights?

The Mist Predictive Analytics and Correlation Engine (PACE) drives the Juniper Mist insights and service level experiences (SLEs). PACE collects the following data.

- Telemetry data from:
 - Juniper wired switches.
 - Edge devices supported by Juniper Mist WAN Assurance.
 - Juniper Mist Edge devices.

- Time to connect data from wireless clients.
- Coverage, roaming, and throughput data from Juniper Mist APs <OR Juniper Access Points>.
- Throughput data for network applications.
- Dwell time and other location data from Bluetooth Low Energy (BLE) tags.

The Mist PACE analyzes and correlates this data to provide you with multiple ways to understand the experiences of users on your network. Use these insights to correct issues, make changes, and ensure a good network experience for your users.

Finding the Insights Dashboard

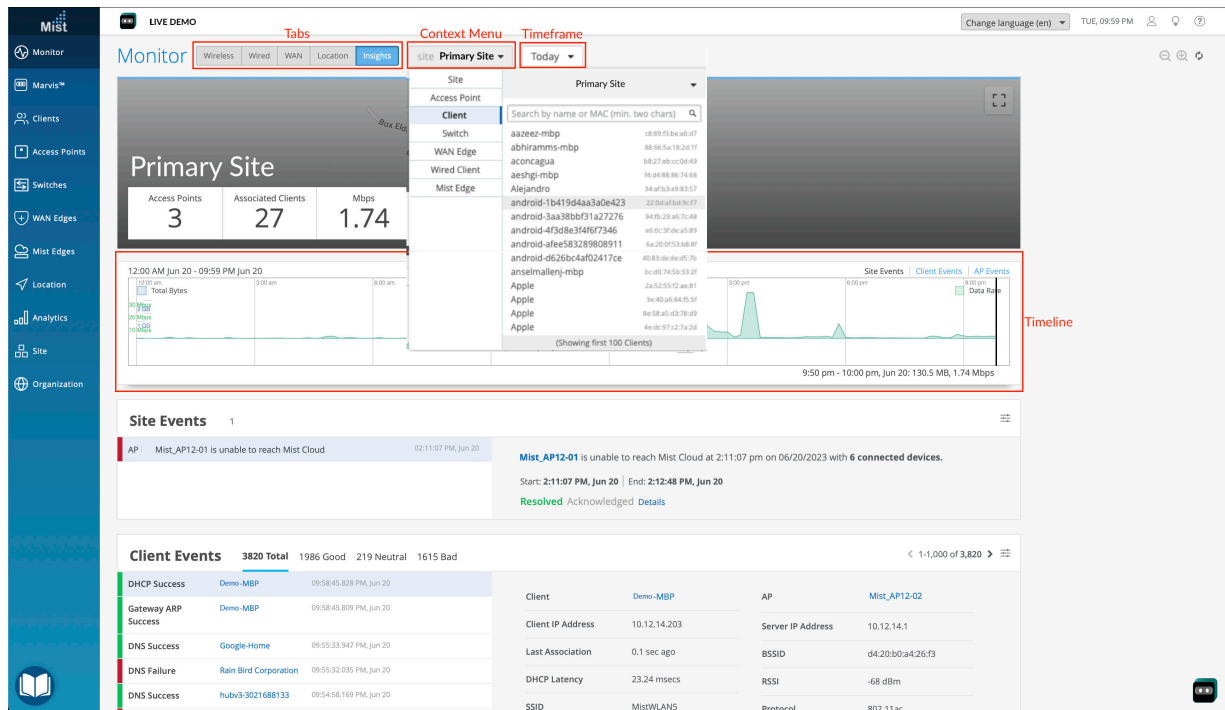
To view the Insights dashboard, select **Monitor** > **Service Levels** from the left menu of the Juniper Mist portal. Then click the **Insights** button at the top of the Monitor page.

Selecting the Context and Time Period

At the top of the Monitor page, select the context, which can be an entire site or a single device or client. In addition, select a time period, such the last 60 minutes, the last 7 days, or a date range.

NOTE: The Monitor page displays data as recent as the past 60 minutes or as far back as the last 7 days. If you purchase a Premium Analytics subscription, you can access up to 3 years' worth of wireless network insights and other data. To access the information available through your Premium Analytics subscription, select **Analytics** > **Premium Analytics** from the left menu of the portal.

Figure 1: Insights Dashboard



Layout of the Monitor Page

The page layout varies, based on the selected context and your active subscriptions. Some of the sections described below might not appear on your screen.

Map Display

Near the top of the Insights dashboard, you'll see a map.

- If you selected a site as the context, the map shows the geographic location of the site.
- If you selected an access point, the map displays the location of the AP within the site floorplan.
- If you selected another context, the map displays that item's location relative to the site or geography that you select from the submenu.

NOTE: If you have not set up a floorplan for your site, the map will be blank. You'll see a grey block with the site name overlaid in white text.

Insights Timeline (Time Range)

Directly below the map, the timeline shows the data rate across the selected time period. You can drag your mouse across the graph to select a time period to zoom in on. The entire page refreshes to show the data for the zoomed-in area of the timeline.

Site Events

In the Site Events block, you can view a list of events that happened within the selected time range at the selected site. Site events apply only to site-assigned APs and RADIUS, DHCP, and DNS servers.

When you select an event from the list, the Insights dashboard shows an information summary about the event. The details link within the event display takes you to the Events page where you can investigate:

- Event actions—Automatic actions, such as sending e-mails or SMS messages, that Mist performed as a result of this event.
- Relevant details—Devices that were impacted, an impact map of the event, and any contributing events that PACE related to this event.

Client Events

In the Client Events block, you can view a list of all events recorded by Mist PACE for the selected site during the selected time frame. These events apply only to wireless clients such as cell phones and laptop computers. When you select an event from the list, Mist shows a summary of the event to the right of the list. If you click the blue text in the summary, you will see details about the client or the AP to which the client was connected at the time of the event.

Juniper access points (APs) have a built-in packet buffer. For certain events such as authorization failures, Mist keeps the buffer information and makes it available as a dynamic packet capture. The Client Events block shows events that have a Dynamic Packet Capture available with a small paper clip

icon next to the event name. A dynamic packet capture can be a very powerful troubleshooting tool. You can download a dynamic packet capture (.pcap file) by clicking the **Download Packet Capture** button in the event summary.

You can filter the Client Events block by clicking the settings button

(



)in the upper-right corner of the block and choosing what to display.

Figure 2: Client Events Filter

Event Filter ✕

All APs **Specific APs**

Event Groups:

- ☒ All Events
- ☒ Connectivity Impacting
- ☒ Fast Roaming
- ☒ Mist Access Assurance (NAC)
- ☐ Connection Setup
- ☒ Roaming Failures
- ☒ Association Failures
- ☒ Client Call Events
- ☒ Connection Failures
- ☒ Slow Roaming
- ☒ Captive Portal Access

Events:

<input checked="" type="checkbox"/> All Positive Events	<input checked="" type="checkbox"/> All Neutral Events	<input checked="" type="checkbox"/> All Negative Events
<input checked="" type="checkbox"/> 11r Association	<input checked="" type="checkbox"/> 802.11 Auth Denied	<input checked="" type="checkbox"/> 11r Auth Failure
<input checked="" type="checkbox"/> 11r FBT Success	<input checked="" type="checkbox"/> AP Deauthentication	<input checked="" type="checkbox"/> 11r FBT Failure
<input checked="" type="checkbox"/> 11r Reassociation	<input type="checkbox"/> Exclude Client Inactivity	<input checked="" type="checkbox"/> 11r Key Lookup Failure
<input checked="" type="checkbox"/> 11r Roam	<input checked="" type="checkbox"/> Client Deauthentication	<input checked="" type="checkbox"/> AirWatch Failure: Not Enrolled
<input checked="" type="checkbox"/> Association	<input checked="" type="checkbox"/> Client Roamed Away	<input checked="" type="checkbox"/> ARP Timed Out
<input checked="" type="checkbox"/> Authentication	<input checked="" type="checkbox"/> DHCP Inform Timed Out	<input checked="" type="checkbox"/> Association Failure
<input checked="" type="checkbox"/> Authorization & Association	<input checked="" type="checkbox"/> Disassociation	<input checked="" type="checkbox"/> Authorization Failure
<input checked="" type="checkbox"/> Authorization & Reassociation	<input type="checkbox"/> Exclude Client Leaving BSS	<input checked="" type="checkbox"/> Bad IP Assigned
<input checked="" type="checkbox"/> Client Joined Call	<input checked="" type="checkbox"/> Local Support Page	<input checked="" type="checkbox"/> Blocked: Policy Lookup Failure
<input checked="" type="checkbox"/> Client Left Call	<input checked="" type="checkbox"/> Portal Redirection Processed	<input checked="" type="checkbox"/> Blocked: Repeated Authorization Failure
<input checked="" type="checkbox"/> DHCP Success	<input checked="" type="checkbox"/> SA Query Timed Out	<input checked="" type="checkbox"/> Blocked: Static DNS Address
<input checked="" type="checkbox"/> DHCPv6 Success		<input checked="" type="checkbox"/> Blocked: Static IP Address
<input checked="" type="checkbox"/> DNS Success		<input checked="" type="checkbox"/> Client Disconnected From Call
<input checked="" type="checkbox"/> Gateway ARP Success		<input checked="" type="checkbox"/> DHCP Denied
<input checked="" type="checkbox"/> MAC Auth Success		<input checked="" type="checkbox"/> DHCP Terminated
<input checked="" type="checkbox"/> NAC Client Access Allowed		<input checked="" type="checkbox"/> DHCP Timed Out
<input checked="" type="checkbox"/> NAC Client Certificate Validation Success		<input checked="" type="checkbox"/> DHCPv6 Denied
<input checked="" type="checkbox"/> NAC IDP Authentication Success		<input checked="" type="checkbox"/> DHCPv6 Terminated
<input checked="" type="checkbox"/> NAC IDP Group Lookup Success		<input checked="" type="checkbox"/> DHCPv6 Timed Out
<input checked="" type="checkbox"/> NAC IDP User Lookup Success		<input checked="" type="checkbox"/> DNS Failure
<input checked="" type="checkbox"/> NAC Server Certificate Validation Success		<input checked="" type="checkbox"/> Excessive ARPing
<input checked="" type="checkbox"/> OKC Reassociation		<input checked="" type="checkbox"/> Gateway ARP Timeout
<input checked="" type="checkbox"/> OKC Roam		<input checked="" type="checkbox"/> Gateway Spoofing
<input checked="" type="checkbox"/> PMKC Association		<input checked="" type="checkbox"/> MAC Auth Failure
<input checked="" type="checkbox"/> PMKC Reassociation		<input checked="" type="checkbox"/> NAC Client Access Denied
<input checked="" type="checkbox"/> Portal Auth Success		<input checked="" type="checkbox"/> NAC Client Certificate Expired
<input checked="" type="checkbox"/> Portal Redirection In Progress		<input checked="" type="checkbox"/> NAC Client Certificate Validation Failure
<input checked="" type="checkbox"/> Reassociation		<input checked="" type="checkbox"/> NAC IDP Admin Config Failure
		<input checked="" type="checkbox"/> NAC IDP Authentication Failure
		<input checked="" type="checkbox"/> NAC IDP Group Lookup Failure
		<input checked="" type="checkbox"/> NAC IDP Lookup Failure
		<input checked="" type="checkbox"/> NAC IDP Unknown
		<input checked="" type="checkbox"/> NAC IDP Unreachable
		<input checked="" type="checkbox"/> NAC IDP User Lookup Failure
		<input checked="" type="checkbox"/> NAC Server Certificate Validation Failure
		<input checked="" type="checkbox"/> OKC Auth Failure
		<input checked="" type="checkbox"/> Portal Auth Failure
		<input checked="" type="checkbox"/> Radius DAS Notify
		<input checked="" type="checkbox"/> SAE Auth Failure

OK **Cancel**

As you can see, Mist provides detailed filtering capabilities for client events.

AP Events

In the AP Events block, you can see a list of AP events that occurred on the selected site during the selected time frame. When you select an event from the list, Mist shows a summary of the event to the right of the list. You can apply similar filters to the AP Events block by clicking the settings button in the upper- right corner of the block.

Applications

Use the Applications block to view a list of the applications in use at the site during the selected time frame. The APs derive the application name primarily from DNS inspection. Mist displays columns of statistics for an application to the right of that application. These statistics are useful for determining how many clients used each application and how much bandwidth the application consumed.

You can click the number of clients to see a list of all the clients that were using the application during the selected time frame. An example of the Applications block is shown below.

Applications 38					
App name	Total Bytes	Percent Bytes	Number of clients	RX Bytes	TX Bytes
Unknown	31.4 GB	56%	44	22.3 GB	9.2 GB
CNN	15.5 GB	28%	4	15.3 GB	241.1 MB
Juniper VPN	3.5 GB	7%	9	2.9 GB	573.9 MB
Yahoo	1.9 GB	4%	2	1.9 GB	20.6 MB
Github	1.7 GB	4%	5	585.9 MB	1.2 GB
Apple	830.5 MB	2%	24	795.1 MB	35.4 MB
Instagram	329.3 MB	1%	5	325.3 MB	4 MB

Network Servers

In the Network Servers block you view see a list of network servers detected in a site. Mist can detect the presence of RADIUS, DHCP, and DNS servers. The data shown to the right of each server in the list can help you spot overused servers and identify those servers with the most failures. This type of data can help you proactively adjust server allocation to enhance the user experience.

Pre-Connection and Post-Connection

You'll see two graphs in the Pre-Connection block, one with DNS latency and the second with DHCP latency. These latency numbers reflect how quickly a wireless client connects to the wireless network, thus affecting the user experience.

You'll see two more graphs in the Post-Connection block, one with the number of connected clients and the second with the number of bytes sent and received. These two graphs show a picture of network attachment and performance, which could also affect user experience.

If you hover over a section of any of the graphs, Juniper Mist updates the display to provide the timestamp and relevant metrics in all the graphs, including the primary time range. With these metrics, you can see trends related to time within the network. You can get a larger view of any of the four Pre-Connection or Post-Connection graphs when you click the expand button

(



) in the upper-left corner of the graph.

NOTE: The following blocks of Mist Insights are not affected by the time range selections. We call these insights Current Values.

Current Site Properties

Use the Current Site Properties block to see information about the selected site at the current time. Along with geographic details, a device count, and the current number of connected clients, you also see a wireless coverage map. The coverage map is based on location information that you provide about the site. See the [Location Services Guide](#) for details about configuring maps and location information.

A heat map is a representation of relative signal strength. The heat map displays an AP as a green circle on the map. The number inside the circle represents the number of wireless clients connected to that particular AP. The dark red background fades to orange and yellow as the distance from the AP grows. This color change represents the change in wireless signal strength associated with each AP.

You can see a larger display of the map when you click the expand button

(



) on the upper-right corner of the wireless signal strength map. You can filter both the small and large maps to show coverage for any available wireless band: 2.4 GHz, 5 GHz, or 6 GHz.

Current WLANs

In the Current WLANs block, you can see wireless LANs configured for the selected site. You can also see the number of APs in the site that host the listed WLAN, the number of clients currently connected to each WLAN, and a summary of traffic and security. When you click on any of the listed WLANs, you can view the WLAN configuration on the configuration page for the WLAN.

Access Points

In the Access Points block, you can see the names of all APs associated with the selected site. Along with the AP name, you can see the connection status, MAC address, uptime, and other information. When you click the name of the AP, the configuration page for that AP appears, where you can view and edit the configuration details.

Clients

You can use the Clients block to see a list of connected clients at the selected site. You can also view offline clients by changing the view to Total. Among the information displayed in the client block is MAC and IP address, device type, and wireless band. When you click on a client name, Juniper Mist takes you to the Client Insights page for that client.

Wired Switches

In the Wired Switches block, you can see a list of EX Series switches associated with the selected site. Summary information includes the switch name, model, IP address, and Junos version information.

3

CHAPTER

Service Level Expectations

[Service-Level Expectations Overview | 24](#)

[Wireless SLEs | 32](#)

[Wired SLEs | 46](#)

[WAN SLEs | 54](#)

[Location SLEs | 59](#)

Service-Level Expectations Overview

SUMMARY

Get familiar with the various elements on the SLE dashboard so that you can use SLEs to assess network performance and address issues as they come up.

IN THIS SECTION

- [What Are Service Level Expectations \(SLEs\)? | 24](#)
- [Finding the SLE Dashboard | 25](#)
- [Selecting the Context and Time Period | 25](#)
- [Using the System Changes Timeline | 26](#)
- [Setting the SLE Thresholds | 27](#)
- [Adjusting the SLE Display Options | 28](#)
- [Understanding the SLE Blocks | 28](#)
- [Sample SLE Block | 30](#)
- [Viewing the Root Cause Analysis Page | 30](#)
- [Summary | 31](#)

What Are Service Level Expectations (SLEs)?

Juniper Mist™ captures, analyzes, correlates, and classifies event and performance data from your network and devices. It then provides you with an assessment of the quality of users' experiences on your network.

Many factors contribute to positive or negative user experiences. Juniper Mist organizes these factors into Service Level Expectations (SLEs). You can set the SLE thresholds to define exactly what "success" means for SLEs such as throughput, capacity, AP health, switch health, and more (as relevant to your network).

When user experiences fail to meet your SLE success thresholds, Juniper Mist identifies the root cause of each poor experience and provides complete details so that you can address the issues.

By skimming the SLE dashboard, you can see at a glance which service levels are low and what types of issues need to be addressed.

The following video gives you a quick, high-level introduction to SLEs.

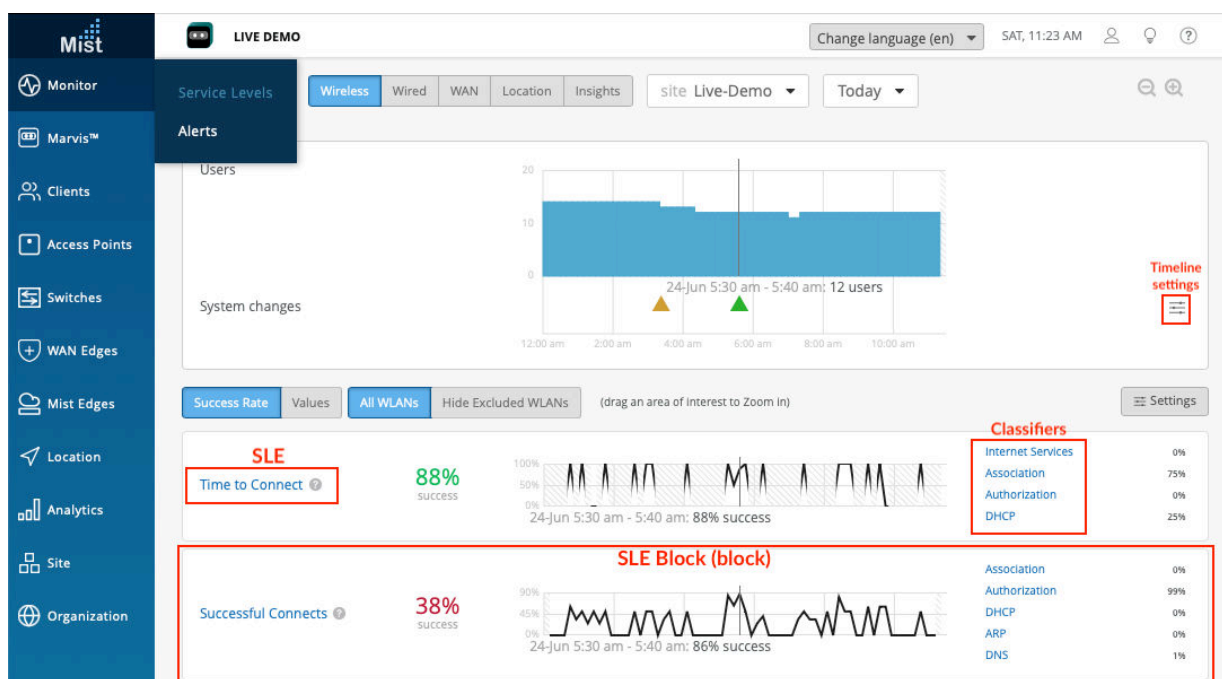


Video:

Finding the SLE Dashboard

To access an SLE dashboard, select **Monitor** > **Service Levels** from the left menu of the Juniper Mist portal. Then use the buttons at the top of the page to select the dashboard that you want to view. Depending on your active subscriptions, these buttons can include Wireless, Wired, WAN, and Location.

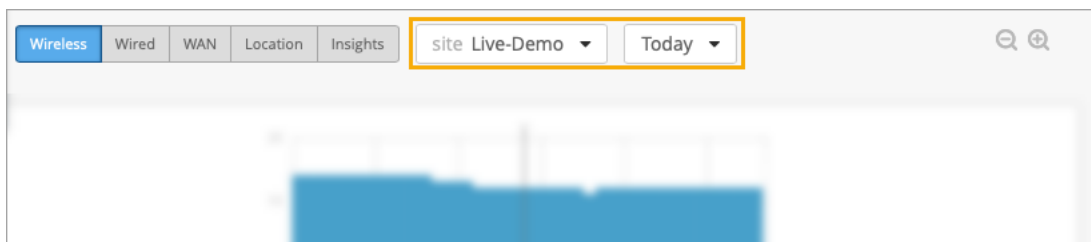
In the following example, you can see the major elements of the SLE dashboard.



NOTE: As you read on, you'll learn about the layout of the dashboard so that you know how to read the display and how to adjust some of the settings. In later topics within this guide, you'll see the definitions of the various SLEs and classifiers for wireless, wired, WAN, and Location.

Selecting the Context and Time Period

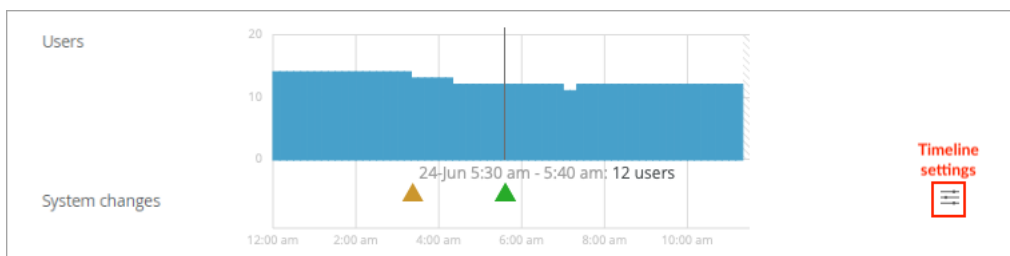
At the top of the Monitor page, select the context, which can be an entire site or a single device or client. In addition, select a time period, such the last 60 minutes, the last 7 days, or a date range.



NOTE: The Monitor page displays data as recent as the past 60 minutes or as far back as the last 7 days. If you purchase a Premium Analytics subscription, you can access up to 3 years' worth of wireless network insights and other data. To access the information available through your Premium Analytics subscription, select **Analytics > Premium Analytics** from the left menu of the portal.

Using the System Changes Timeline

When investigating issues, your first question might be, "Did anything change on the network?" With this timeline, you can see at a glance if any system changes occurred and how many users or clients were active at the time.



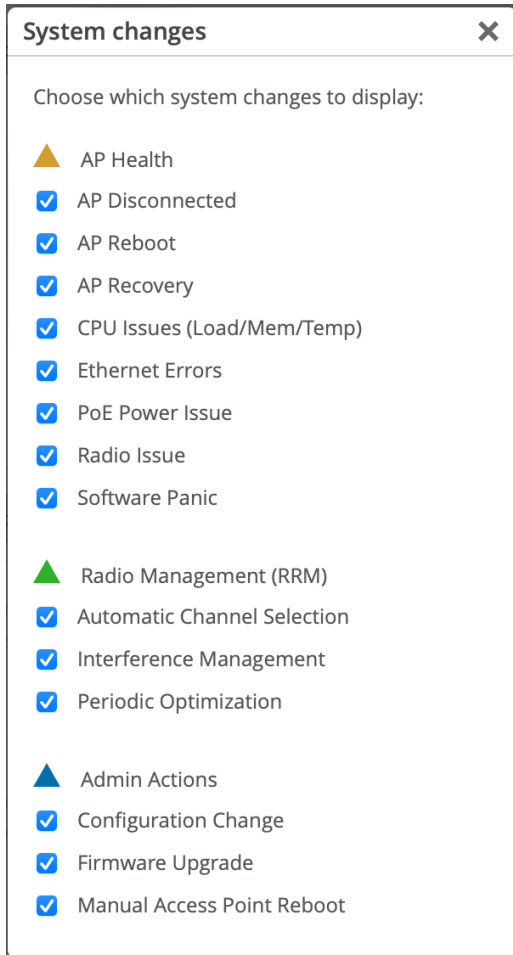
The triangles below the timeline represent various types of system changes:

- Yellow triangle—AP Health
- Green triangle—Radio Management (RRM)
- Blue triangle—Admin Actions

You can adjust the timeline settings to specify the types of changes to include. To get started, click the timeline settings button:



In the System Changes window, select or deselect check boxes for each event that you want to include or exclude.



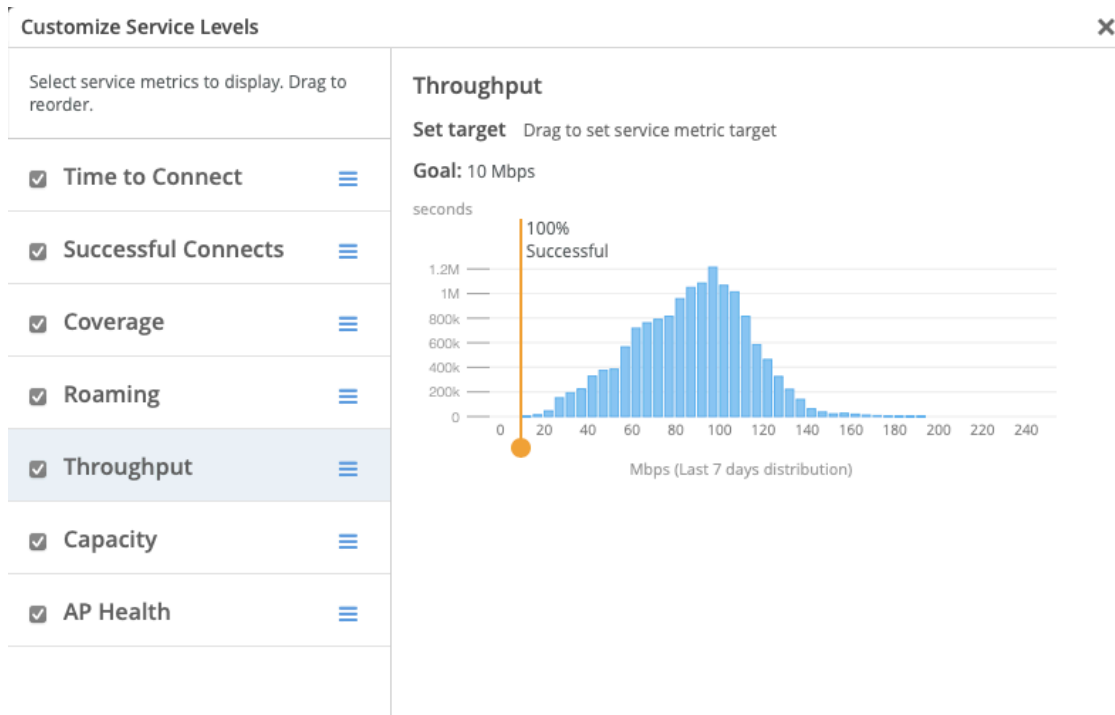
Setting the SLE Thresholds

Each SLE has a success threshold. For the Time to Connect SLE, for example, you might set a threshold of 2 seconds. This means that you consider your network successful when users can send and receive data over the Internet within 2 seconds of attempting to associate with an access point.

To view or modify the SLE thresholds, you can click the **Settings** button on the right side of the SLE dashboard.

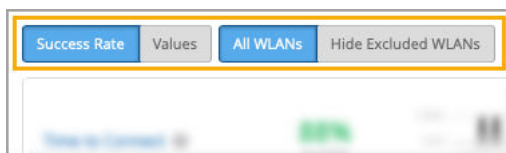


In the Customize Service Levels window, you can modify the thresholds as needed to ensure that the SLE settings meet your goals for your network.



NOTE: This example shows the wireless SLEs. Depending on the dashboard that you're viewing, you'll see different SLEs in this window.

Adjusting the SLE Display Options



You can adjust the SLE dashboard display as follows:

- Show success rates or values.
- Include all WLANs or hide excluded WLANs.

Understanding the SLE Blocks

Each SLE is represented by a separate block (sub-section) on the dashboard.

In each block, you'll see:

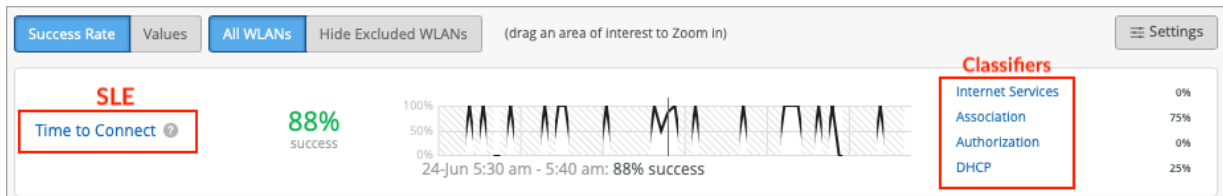
- **Overall Service Level.** On the left side of each SLE block, you'll see the overall service level for the selected site and time period.
 - Click **Success Rate** to see the *percentage* of user experiences that met the SLE success threshold.
 - Click **Values** to see the *number* of user experiences that met the SLE success threshold.
- **Timeline.** In the middle of each SLE block, you can explore the timeline. As your mouse moves across the timeline, information appears under it.
 - Click **Success Rate** to see the *percentage* of successful user experiences at the selected point in time.
 - Click **Values** to see the *number* of successful user experiences at the selected point in time.
- **Classifiers.** On the right side of each SLE block, you see the *classifiers* for the user experiences that didn't meet the SLE success threshold. Juniper Mist attributes each unsuccessful user experience to one classifier. Together, the classifiers give you a high-level root cause analysis of the unsuccessful user experiences.
 - Click **Success Rate** to see the *percentage* of unsuccessful user experiences that were caused by each classifier.

NOTE: Together, these individual percentages total 100 percent of the unsuccessful user experiences.

- Click **Values** to see the *number* of unsuccessful user experiences that were caused by each classifier.

NOTE: Together, these individual values represent the total number of unsuccessful user experiences.

Sample SLE Block



In this example, Success Rate button is selected, so you see percentages instead of values.

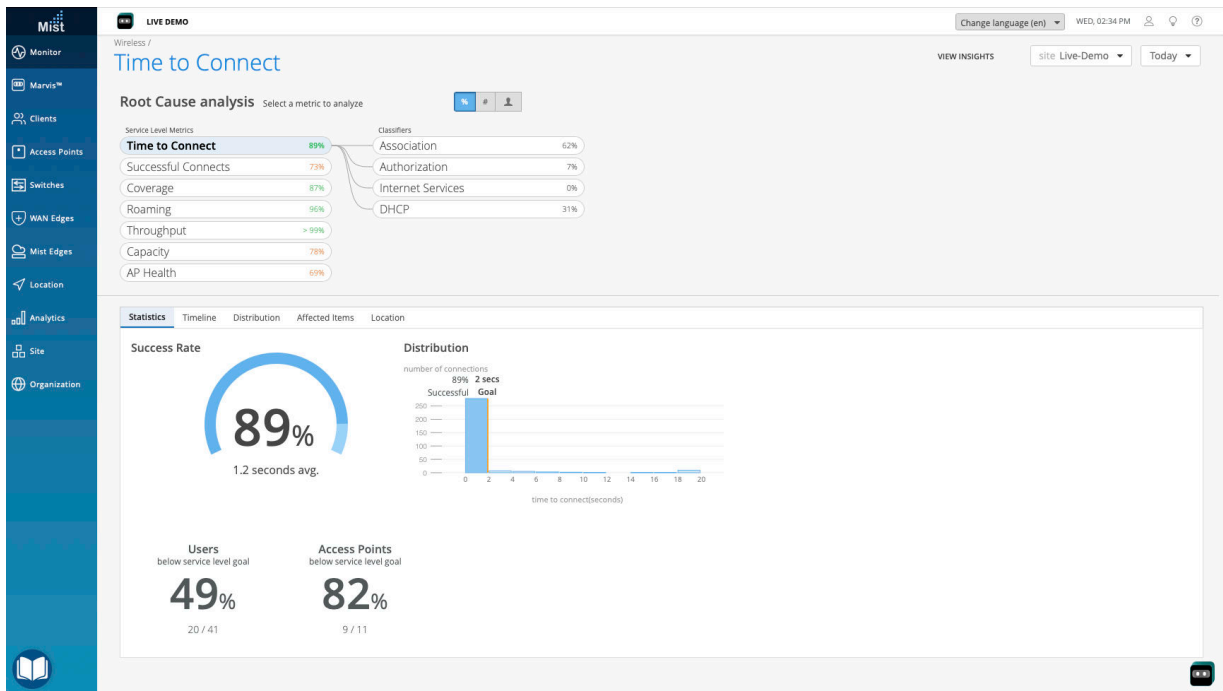
- On the left, you see that the overall success rate for the selected site and time period was 88 percent.
- In the middle, the timeline caption shows that the mouse is hovering over 24-Jun 5:30 am - 5:40 am. At that point, the success rate was 88 percent.

On the right, you see that 75 percent of the SLE-lowering issues occurred in the Association process and 25 percent occurred in the DHCP process. Together, these classifiers account for 100 percent of the user experiences that failed to meet the threshold. The other classifiers show 0 percent, meaning that they did not have any impact on this SLE.

Viewing the Root Cause Analysis Page

From the dashboard, you can click any SLE or classifier to go to the Root Cause Analysis page.

This example shows the Root Cause Analysis page for the wireless Time to Connect SLE.



- At the top of the page, you see the data for all classifiers and their sub-classifiers (if applicable).
- In the lower part of the page, you see additional details about the selected item. Depending on the classifier, you might see signal strength information, a list of affect devices and clients, or other information. These details help you to understand the scope of the issues.

Summary

Now you're familiar with the layout of the SLE dashboard, some of the settings that you can change, and how to go to the Root Cause Analysis page. Start exploring the SLEs and classifiers for your network.

- ["Wireless SLEs" on page 32](#)
- ["Wired SLEs" on page 46](#)
- ["WAN SLEs" on page 54](#)
- ["Location SLEs" on page 59](#)

Wireless SLEs

SUMMARY

Use the wireless service-level experience (SLE) dashboard to assess the service levels for user-impacting factors such as throughput, signal strength, roaming, and more.

IN THIS SECTION

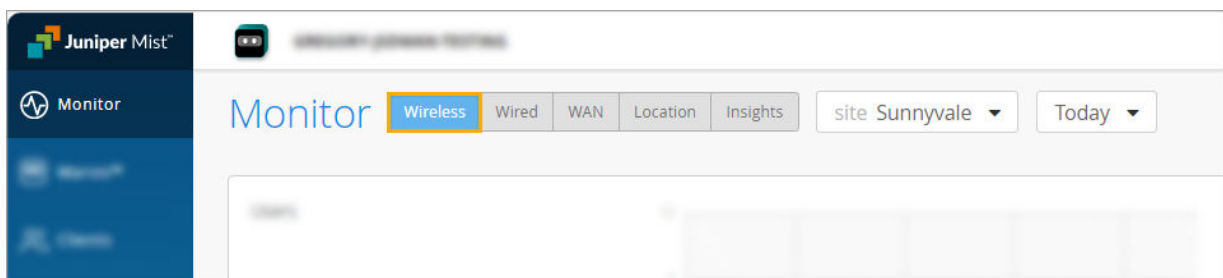
- [Time to Connect SLE | 33](#)
- [Wireless Successful Connects SLE | 35](#)
- [Coverage SLE | 37](#)
- [Roaming SLE | 39](#)
- [Wireless Throughput SLE | 41](#)
- [Capacity SLE | 43](#)
- [AP Health SLE | 44](#)

Finding the Wireless SLE Dashboard

The Wireless SLE Dashboard is part of the Monitor page, which appears automatically after you log in to the Juniper Mist™ portal.

From other pages of the portal, you can open the Monitor page by selecting **Monitor > Service Levels** from the left menu.

On the Monitor page, click the **Wireless** button to view the Wireless SLE Dashboard.



Wireless Service Level Expectations Video Overview



Video:

Using the Wireless SLE Dashboard

For a general introduction to the SLE dashboard, see ["Service-Level Expectations Overview" on page 24](#).

For help interpreting the wireless SLEs and classifiers, go to these topics:

- ["Time to Connect SLE" on page 33](#)
- ["Wireless Successful Connects SLE" on page 35](#)
- ["Coverage SLE" on page 37](#)
- ["Roaming SLE" on page 39](#)
- ["Wireless Throughput SLE" on page 41](#)
- ["Capacity SLE" on page 43](#)
- ["AP Health SLE" on page 44](#)

Time to Connect SLE

SUMMARY

Use the Time to Connect SLE to assess your users' experience connecting to the Internet through your wireless network.

IN THIS SECTION

- [What Does the Time to Connect SLE Measure? | 33](#)
- [Classifiers for Excess Time to Connect | 34](#)
- [Example | 34](#)

Time to Connect is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Time to Connect SLE Measure?

Time to Connect is the number of seconds that elapse between the point when a client sends an association packet and the moment when the client can successfully move data.

You can click the **Settings** button to set the number of seconds to use as the success threshold for this SLE.

Classifiers for Excess Time to Connect

When the Time to Connect threshold is not met, Juniper Mist classifies the issues as follows.

- **Internet Services**—The time to access external networks was more than 2 sigma from the moving average for this site.
- **Authorization**—The time to go past the authentication state was more than 2 sigma from the average authentication latency for this site.
- **Association**—The time to go past the association state was more than 2 sigma from the average association latency for this site.
- **DHCP**—The time to connect to Dynamic Host Configuration Protocol (DHCP) was more than 2 sigma from the average time for fully completed successful connections for this site.

Sub-Classifiers for DHCP:

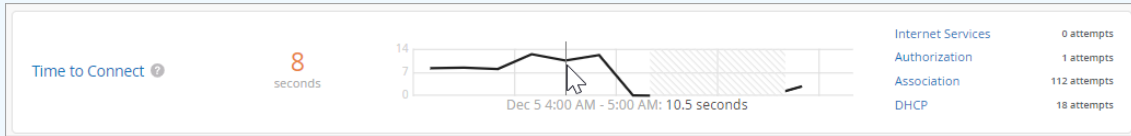
- Stuck
- Nack
- Unresponsive

Example



- **Success Rate**—On the left, you see that the Time to Connect threshold was met by only 51 percent of the user connections.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. In this example, there was a 14 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the connections that failed to meet this SLE. Of these connections, 86 percent had trouble with Association and 14 percent had trouble with DHCP. No issues (0 percent) were attributed to the other classifiers.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. In this example, the left side of the screen shows the average number of seconds to connect. The timeline shows the number of seconds for the selected time. The right side of the screen shows the number of connection attempts that were impacted by each classifier.



Wireless Successful Connects SLE

SUMMARY

Use the Wireless Successful Connects SLE to assess your users' experiences connecting to your wireless network.

IN THIS SECTION

- [What Does the Wireless Successful Connects SLE Measure? | 35](#)
- [Classifiers for Unsuccessful Connections | 35](#)
- [Example | 36](#)

Successful Connects is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Wireless Successful Connects SLE Measure?

Juniper Mist tracks the success or failure of authorization, association, DHCP, ARP, and DNS attempts. These connection attempts include initially connecting to the network, roaming from one AP to another, and on-going connectivity.

You don't need to set up a threshold for this SLE. It's assumed that you want 100 percent successful connects.

Classifiers for Unsuccessful Connections

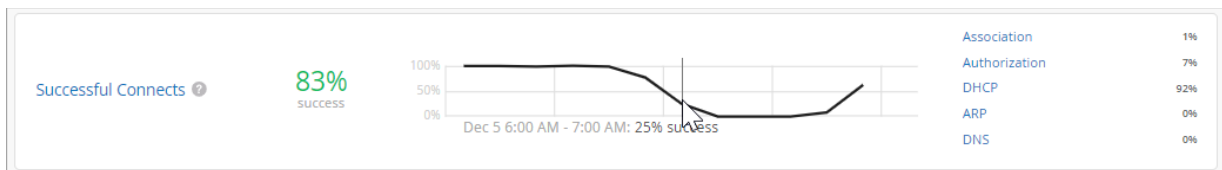
When connection attempts fail, Juniper Mist classifies the issues as follows.

- **Association**—The connection failed during the association process.
- **Authorization**—The connection failed during the authorization process.
- **DHCP**—The connection failed during the DHCP process.

The DHCP classifier has four sub-classifiers:

- Renew Unresponsive
- Nack
- Incomplete
- Discover Unresponsive
- **ARP**—The client experienced one of these problems:
 - ARP failure for the default gateway during the initial connection
 - ARP gateway failures after the initial connection or roam
- **DNS**—The client experienced DNS failures during or after the connection process.

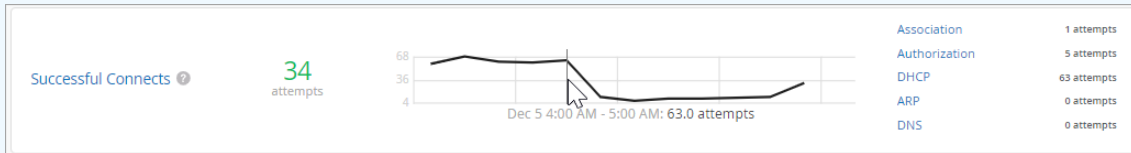
Example



- **Success Rate**—On the left, you see that only 83 percent of connection attempts succeeded.
- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was a 25 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the failed attempts. The vast majority of the issues (92 percent) happened during the DHCP process. Another 7 percent failed during authorization, and 1 percent failed during association. No issues (0 percent) were attributed to the other classifiers.

NOTE: You can click the **Values** button to show numbers instead of the success rate and the classifier percentages. In this example, the left side of the screen shows the number of successful connects during the specified time period. The timeline shows the number of attempts at each

point in time. The right side of the screen shows the number of failed connection attempts for each classifier.



Coverage SLE

SUMMARY

Use the Coverage SLE to assess your users' experiences with signal strength.

IN THIS SECTION

- [What Does the Coverage SLE Measure? | 37](#)
- [Classifiers for Poor Coverage | 37](#)
- [Example | 38](#)

Coverage is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Coverage SLE Measure?

Juniper Mist tracks active clients' Received Signal Strength Indicator (RSSI), as measured by the access point.

You can click the **Settings** button to set the RSSI level that you want to use as the success threshold for this SLE.

Classifiers for Poor Coverage

When the RSSI threshold is not met, Juniper Mist classifies the issues as follows.

- **Asymmetry Uplink**—Clients received a weak signal due to asymmetric uplink strength between the AP and the client device. (Uplink traffic is the traffic going from the client to the AP, and then to the Internet.) Asymmetry can occur for various reasons, such as clients being too far from the AP.

- **Asymmetry Downlink**—Clients received a weak signal due to asymmetric downlink transmission strength between the AP and a client device. (The traffic going from the AP to the client is called downlink traffic.)
- **Weak Signal**—Clients received a weak signal due to other factors.

Example



- **Success Rate**—On the left, you see that the threshold was met 94 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was a 96 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the user experiences that were impacted by poor coverage. The majority of the issues (63 percent) were due to Asymmetry Uplink. Another 36 percent were due to Weak Signal, and 1 percent were due to Asymmetry Downlink.

NOTE: You can click the **Values** button to show numbers instead of the success rate and the classifier percentages. In this example, the left side of the screen shows the average RSSI during the specified time period. The timeline shows the average RSSI at each point in time. The right side of the screen shows the number of user minutes that were impacted by each classifier.



Roaming SLE

SUMMARY

Use the Roaming SLE to track successful and unsuccessful roams between access points.

IN THIS SECTION

- [What Does the Roaming SLE Measure? | 39](#)
- [Classifiers for Poor Roaming | 39](#)
- [Example | 40](#)

Roaming is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Roaming SLE Measure?

Juniper Mist tracks the percentage of successful roams between access points and assigns a quality score from 1 to 5. A score of 1 indicates excellent roaming, and a score of 5 indicates poor roaming.

You can click the **Settings** button to set the roaming score to use as the success threshold for this SLE.

Classifiers for Poor Roaming

When the roaming threshold is not met, Juniper Mist classifies the issues as follows.

- **Latency**—Roaming time was excessive.

Latency has different sub-classifiers for different roaming options:

- **Slow 11r Roams**—This classifier applies to fast roaming as defined by [802.11r](#). The roaming time exceeded 400 minutes.
- **Slow Standard Roams**—This classifier applies to standard roaming. The roaming time exceeded 2 seconds.
- **Slow OKC Roams**—This classifier applies to clients using RADIUS-based authentication with Opportunistic Key Caching (OKC). The roaming time exceeded 2 seconds.
- **Stability**—This classifier tracks the consistency of AP choice and 11r usage during client roams. Juniper Mist assigns this classifier if a user capable of fast roaming on a fast-roaming enabled SSID experiences slow roaming for more than 2 seconds. This classifier contains one sub-classifier: **Failed to fast Roam**.

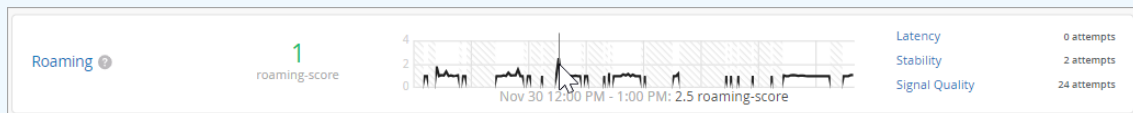
- **Signal Quality**—This classifier tracks the RSSI of clients during a roaming event.
- **Suboptimal Roam**—This sub-classifier tracks when clients roam to an AP:
 - With more than 6 dBm decrease in RSSI compared to the client's RSSI in the previous AP
 - If the RSSI in the new connection is worse than the configured coverage SLE threshold. Note that the default coverage SLE threshold is 72 dBm.
- **Sticky Client**—This sub-classifier tracks the events when a client remains connected to an AP even when more roaming options are available to improve the RSSI by more than 6 dBm.

Example



- **Success Rate**—On the left, you see that 99 percent of roaming attempts met the threshold.
- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was a 100 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the roaming attempts that did not meet the threshold. The vast majority of the issues (92 percent) were attributed to Signal Quality. The other 8 percent were attributed to Stability. No issues (0 percent) were attributed to the other classifier.

NOTE: You can click the **Values** button to show numbers instead of the success rate and the classifier percentages. In this example, the left side of the screen shows the average roaming score during the specified time period. The timeline shows the average roaming score at each point in time. The right side of the screen shows the number of roaming attempts that were impacted by each classifier.



Wireless Throughput SLE

SUMMARY

Use the Throughput SLE to assess users' experiences with throughput on your wireless network.

IN THIS SECTION

- [What Does the Wireless Throughput SLE Measure? | 41](#)
- [Classifiers for Low Throughput | 41](#)
- [Example | 42](#)

Throughput is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Wireless Throughput SLE Measure?

Juniper Mist calculates the estimated throughput on a per-client basis for the entire site. The estimator considers effects such as AP bandwidth, load, interference events, the type of wireless device, signal strength, and wired bandwidth, to arrive at the probabilistic throughput.

You can click the **Settings** button to set the success threshold for this SLE.

Classifiers for Low Throughput

When the throughput threshold is not met, Juniper Mist classifies the issues as follows.

- **Device Capability**—Low throughput is primarily due to the device capability.
- **Capacity**—Low throughput is due to one of these factors:
 - Load on the associated AP.
 - Wireless or non-wireless interference on the channel.

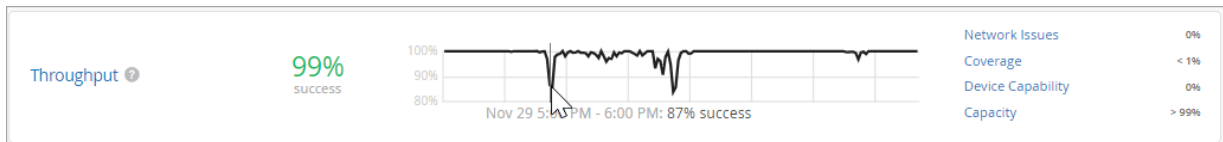
The capacity classifier has four sub-classifiers:

- High Bandwidth Utilization
- Non Wi-Fi Interference
- Excessive Client Load
- Wi-Fi Interference

You can use these sub-classifiers to analyze users and APs below the SLE goal, the timeline of failures and system changes, and the distribution of failures. You can also analyze related network processes that these sub-classifiers can influence.

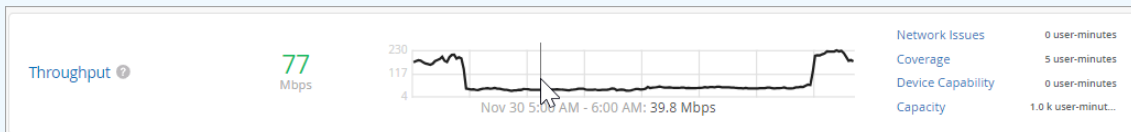
- **Coverage**—Low throughput is primarily due to the client's weak signal strength.
- **Network Issues**—Low throughput is primarily due to the capacity of the wired network.

Example



- **Success Rate**—On the left, you see that the threshold was met 99 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was an 87 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the user minutes that did not meet the threshold. More than 99 percent of these issues were due to Capacity. Less than 1 percent were due to Coverage.

NOTE: You can click the **Values** button to show numbers instead of the success rate and the classifier percentages. In this example, the left side of the screen shows the average throughput during the specified time period. The timeline shows the average throughput at each point in time. The right side of the screen shows the number of user minutes that were impacted by each classifier.



Capacity SLE

SUMMARY

Use the Capacity SLE to track user experiences with RF channel capacity (bandwidth) on your wireless network.

IN THIS SECTION

- [What Does the Capacity SLE Measure? | 43](#)
- [Classifiers for Low Capacity | 43](#)
- [Example | 43](#)

Capacity is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Capacity SLE Measure?

Juniper Mist monitors the percentage of the total RF channel capacity that is available to clients.

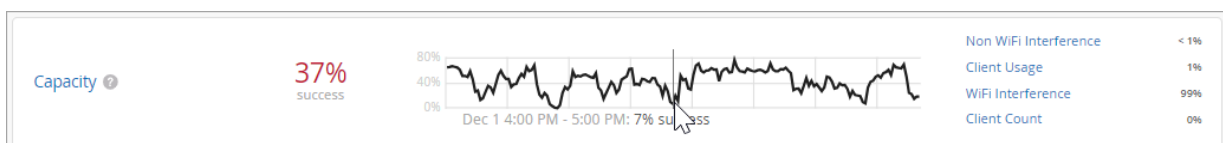
You can click the **Settings** button to set the success threshold for this SLE. For example, you might want 20 percent of the RF channel capacity (bandwidth) to be available to clients at any time.

Classifiers for Low Capacity

When the capacity threshold is not met, Juniper Mist classifies the issues as follows.

- **Wi-Fi interference**—Low capacity is due to wireless interference.
- **Non-Wi-Fi interference**—Low capacity is due to non-wireless interference.
- **Client Count**—Low capacity is due to a high number of attached clients.
- **Client Usage**—Low capacity is due to a high client load.

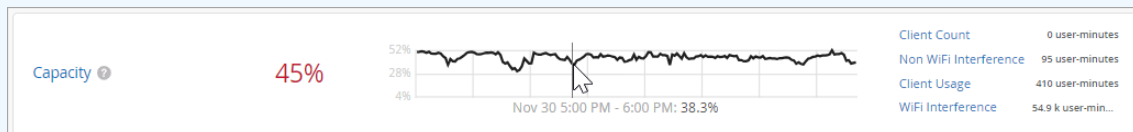
Example



- **Success Rate**—On the left, you see that the threshold was met only 37 percent of the time.

- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was a 7 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the user minutes that did not meet the threshold. Almost all these issues (99 percent) were due to Wi-Fi interference. The remaining issues were due to other interference and client usage. No issues (0 percent) were attributed to Client Count.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. With the Capacity SLE, you see the average percentage of bandwidth that was available during the specified time period. The timeline shows the percentage of total bandwidth that was available at each point in time. The right side of the screen shows the number of user minutes that were impacted by each classifier.



AP Health SLE

SUMMARY

Use the AP Health SLE to assess your users' experience with AP availability.

IN THIS SECTION

- [What Does the AP Health SLE Measure? | 45](#)
- [Classifiers for Poor AP Health | 45](#)
- [Example | 45](#)

AP Health is one of the wireless Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the AP Health SLE Measure?

Juniper Mist tracks the percentage of time the APs are operational without rebooting or losing connectivity to the cloud.

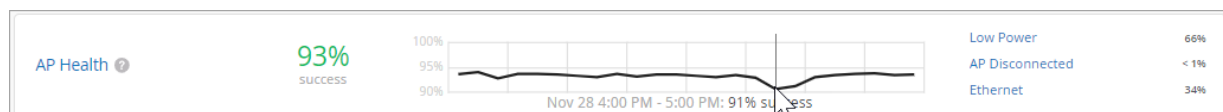
You don't need to set up a quality threshold for this SLE. It's assumed that you do not want *any* AP health issues.

Classifiers for Poor AP Health

When AP Health is poor, Juniper Mist classifies the issues as follows.

- **Low Power**—An AP received insufficient power from its Power over Ethernet (PoE) connection.
- **AP Disconnected**—One of these conditions occurred:
 - Switch Down—Multiple APs that were connected to the same switch lost cloud connectivity.
 - Site Down—All the APs on the site were unreachable.
 - AP Unreachable—An AP lost cloud connectivity.
 - AP Reboot—An AP rebooted.
- **Ethernet**—One of these conditions occurred:
 - Speed Mismatch—Juniper Mist detected a speed or duplex mismatch between an upstream device and an AP.
 - Ethernet Errors—Juniper Mist detected cyclic redundancy check (CRC) errors on the Ethernet interface of the AP.

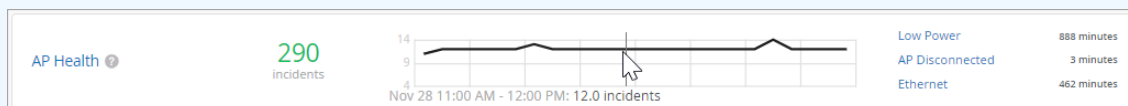
Example



- **Success Rate**—On the left, you see that the APs were available 93 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore the success rate over time. In this example, there was a 91 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.

- **Classifiers**—On the right, you see a high-level root cause analysis for the user minutes that did not meet the threshold. Most of these issues (66 percent) were due to low power. Another 34 percent of the issues were due to Ethernet problems. Less than 1 percent were classified as AP Disconnected.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see the number of incidents during this time period. The timeline shows the number of incidents at each point in time. On the right, you see the number of user minutes that APs were impacted by each classifier.



Wired SLEs

SUMMARY

Use the wired service-level experience (SLE) dashboard to assess the service levels for user-impacting factors such as throughput, connectivity, and switch health.

IN THIS SECTION

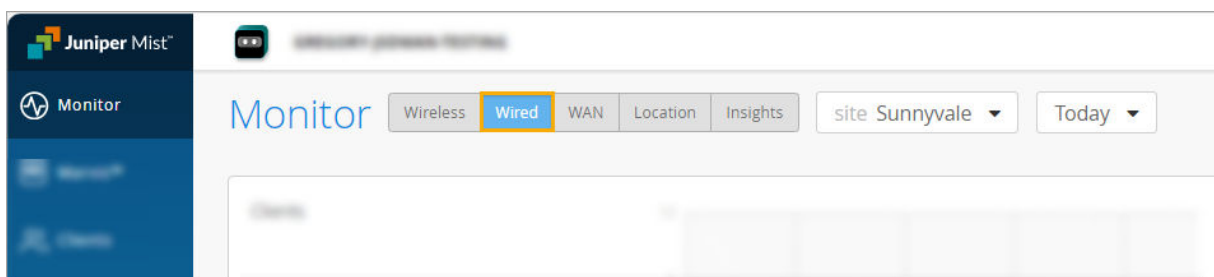
- [Wired Throughput SLE | 47](#)
- [Switch Health SLE | 50](#)
- [Wired Successful Connect SLE | 52](#)

Finding the Wired SLE Dashboard

The Wired SLE Dashboard is part of the Monitor page, which appears automatically after you log in to the Juniper Mist™ portal.

From other pages of the portal, you can open the Monitor page by selecting **Monitor > Service Levels** from the left menu.

On the Monitor page, click the **Wired** button to view the Wired SLE Dashboard.



Wired Assurance: Day 2 - Wired Service Level Expectations (SLEs) Video Overview



Video:

Using the Wired SLE Dashboard

For a general introduction to the SLE dashboard, see ["Service-Level Expectations Overview" on page 24](#).

For help interpreting the wired SLEs and classifiers, go to these topics:

- ["Wired Throughput SLE" on page 47](#)
- ["Switch Health SLE" on page 50](#)
- ["Wired Successful Connect SLE" on page 52](#)

Wired Throughput SLE

SUMMARY

Use the Wired Throughput SLE to assess users' experiences with throughput on your wired network.

IN THIS SECTION

- [What Does the Wired Throughput SLE Measure? | 48](#)

●	Classifiers for Low Throughput 48
●	Example 49

Throughput is one of the wired Service-Level Expectations (SLEs) that you can track on the wired SLE dashboard in the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Wired Throughput SLE Measure?

Juniper Mist measures the available bandwidth on your network. This SLE can help you to determine if you need more wired bandwidth on your site.

You can click the **Settings** button to set the success threshold for this SLE.

Classifiers for Low Throughput

When the throughput threshold is not met, Juniper Mist classifies the issues as follows.

- **Congestion Uplink**—The SLE dashboard shows high congestion uplink when:
 - One of the neighbors is a switch or a router (known through LLDP).
 - The port is a Spanning Tree Protocol (STP) root port.
 - The uplink port has a higher number of transmitted and received packets compared to the other ports.
 - Aggregated Links. Congestion can also be caused by aggregated Ethernet links and module ports.
- **Congestion**—This classifier measures the number of output drops. When packets come into a switch interface, they are placed in an input queue (buffer). When the buffer becomes full, it will start to drop packets (Tx Drops). We use a formula that takes into account the following ratios to determine if there is a 'bad user minute' due to congestion:
 - Tx Drops to Tx Packets—Total transmitted bytes dropped to total packets transmitted.
 - Txbps to Link speed—Total bytes transmitted per second to link speed.
 - RxSpeed to Link speed—Total bytes received per second to link speed.
- **Storm Control**—Storm control allows the device to monitor traffic levels and drop broadcast, unknown unicast, and multicast packets when they exceed a set threshold or traffic level. This threshold is known as a storm control level or storm control bandwidth. The default storm control

level is 80 percent of the combined broadcast, multicast, and unknown unicast traffic on all Layer 2 interfaces of Juniper switches. Storm control helps prevent traffic storms, but it can also potentially throttle applications or client devices. This classifier identifies these conditions and helps users proactively mitigate throughput issues.

- **Network**—You can use this classifier to monitor user minutes when the throughput is lower than expected due to uplink capacity limitations. It identifies issues based on the round-trip time (RTT) value of packets sent from the switch to the Mist cloud. The Network classifier has two sub-classifiers that help you identify these issues:
 - **Latency**—Displays user minutes affected by latency. The latency value is calculated based on the average value of RTT over a period of time.
 - **Jitter**—Displays user minutes affected by jitter. The jitter value is calculated by comparing the standard deviation of RTT within a small period (last 5 or 10 minutes) with the overall deviation of RTT over a longer period (day or week). You can view this information for a particular switch or site.
- **Interface Anomalies**—The details for interface anomalies are all obtained from the switch. The Interface Anomalies classifier contains three sub-classifiers: MTU Mismatch, Cable Issues, and Negotiation Failed.
 - **MTU Mismatch**—As an administrator, you can set an MTU value for each interface. The default value for Gigabit Ethernet interfaces is 1514 . To support jumbo frames, you must configure an MTU value of 9216, which is the upper limit for jumbo frames on a routed virtual LAN (VLAN) interface. It's important to ensure that the MTU value is consistent along the packet's path, as any MTU mismatch will result in discarded or fragmented packets. In Juniper Networks switches, you can check for MTU mismatches in the **MTU Errors** and **Input Errors** sections of the show interface extensive command output. Each input error or MTU error contributes to a "bad user minute" under MTU mismatch.
 - **Cable Issues**—This sub-classifier shows the user minutes affected by faulty cables in the network.
 - **Negotiation Failed**—Latency on ports can happen due to autonegotiation failure, duplex conflicts, or user misconfiguration of device settings. Moreover, older devices may fail to achieve maximum speed and could operate at a slower link speed of 100 Mbps. This sub-classifier identifies and helps mitigate instances of bad user time caused by these issues.

Example



- **Success Rate**—On the left, you see that the threshold was met 99 percent of the time.

- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the unsuccessful user experiences. Most (80 percent) were due to Congestion, 19 percent were due to Interface Anomalies, 1 percent to Storm Control, and less than 1 percent to Congestion Uplink.

NOTE: If you click the **Values** button above the SLE blocks, you'll see numbers instead of percentages for the success rate, timeline, and classifiers.

Switch Health SLE

SUMMARY

Use the Switch Health SLE to assess switch performance and to identify user-impacting issues with switch reachability, memory, CPU, and more.

IN THIS SECTION

- [What Does the Switch Health SLE Measure? | 50](#)
- [Classifiers for Poor Switch Health | 51](#)
- [Example | 51](#)

Switch Health is one of the Wired Service-Level Expectations (SLEs) that you can track on the Wired SLE dashboard in the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Switch Health SLE Measure?

Juniper Mist™ monitors your switches' operating temperatures, power consumption, CPU, and memory usage. Monitoring switch health is crucial because issues such as high CPU usage can directly impact connected clients. For instance, if CPU utilization spikes to 100 percent, the connected APs may lose connectivity, affecting the clients' experience.

Juniper Mist assigns a quality score from 1 to 5. A score of 1 indicates excellent health, and a score of 5 indicates poor health.

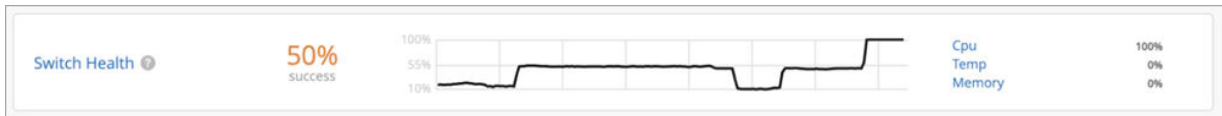
You can click the **Settings** button to set the score to use as the success threshold for this SLE.

Classifiers for Poor Switch Health

When the Switch Health threshold is not met, Juniper Mist classifies the issues as follows.

- **Switch Unreachable**—The switch can't be accessed.
- **Memory**—The memory utilization is above 80 percent.
- **CPU**—The CPU usage of the switch is above 90 percent.
- **Temp**—The operating temperature of the switch is outside the prescribed threshold range, going either above the maximum limit or below the minimum requirement.
- **Power**—The switch is consuming over 90 percent of the available power.

Example



- **Success Rate**—On the left, you see that the threshold was met 50 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the switch health issues. All (100 percent) were attributed to CPU. No issues (0 percent) were attributed to the other classifiers.

NOTE: If you click the **Values** button above the SLE blocks, you'll see numbers instead of percentages for the success rate, timeline, and classifiers.

Wired Successful Connect SLE

SUMMARY

Use the Wired Successful Connect SLE to assess clients' experiences connecting to your wired network.

IN THIS SECTION

- [What Does the Wired Successful Connect SLE Measure? | 52](#)
- [Classifiers for Unsuccessful Connection Attempts | 52](#)
- [Example | 53](#)

Successful Connect is one of the Wired Service-Level Expectations (SLEs) that you can track on the Wired SLE dashboard in the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Wired Successful Connect SLE Measure?

Juniper Mist monitors client connection attempts and identifies failures. This SLE helps you to assess the impact of these failures and to identify the issues to address.

You do not need to set a threshold for this SLE. Juniper Mist assumes a success threshold of 100 percent, meaning that you want all connection attempts to succeed.

Classifiers for Unsuccessful Connection Attempts

When connection attempts are unsuccessful, Juniper Mist classifies the issues as follows.

- **Authentication**—Each time a client authenticates, a client event is generated. These could either be successful or failed events. This classifier helps you identify issues that caused authentication failures. Here's a list of possible reasons for an 802.1X authentication failure:
 - If a single switch port fails to authenticate, it could be due to a user error or misconfigured port.
 - If all switch ports fail to authenticate, it could be because:
 - The switch is not added as a NAS client in the RADIUS server.
 - A routing issue exists between the switch and the RADIUS server.
 - The RADIUS server is down.

- If all switch ports on all the switches fail to authenticate, it could indicate a temporary failure with the RADIUS server at that specific moment.
- If a specific type of device, such as a Windows device, fails to authenticate, it may suggest an issue related to certifications.
- **DHCP**—Dynamic Host Configuration Protocol (DHCP) snooping enables the switch to examine the DHCP packets and keep track of the IP-MAC address binding in the snooping table. This classifier adds a failure event every time a client connects to a network and fails to reach the 'bound' state within a minute.

NOTE:

The SLE dashboard shows DHCP failures only for those switches that have DHCP snooping configured.

Example



- **Success Rate**—On the left, you see that 49 percent of connection attempts were successful.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the unsuccessful connection attempts. All (100 percent) occurred during the Authentication process. None occurred during the DHCP process.

NOTE: If you click the **Values** button above the SLE blocks, you'll see numbers instead of percentages for the success rate, timeline, and classifiers.

WAN SLEs

SUMMARY

Use the WAN service-level experience (SLE) dashboard to assess the service levels for user-impacting factors such as WAN Edge health, WAN link health, and application health.

IN THIS SECTION

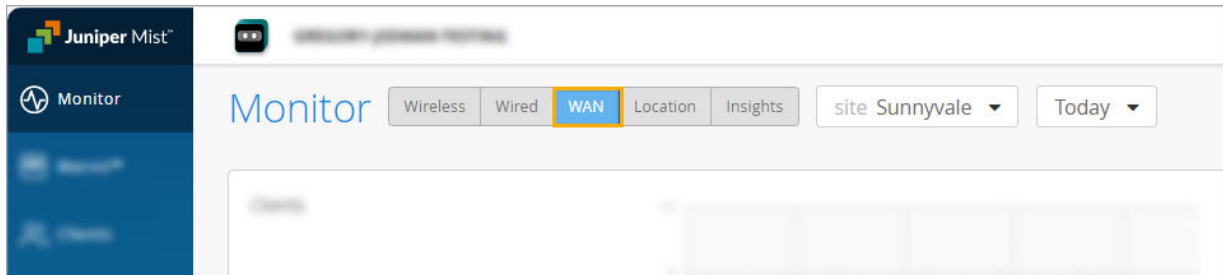
- [WAN Edge Health SLE | 55](#)
- [WAN Link Health SLE | 57](#)

Finding the WAN SLE Dashboard

The WAN SLE Dashboard is part of the Monitor page, which appears automatically after you log in to the Juniper Mist™ portal.

From other pages of the portal, you can open the Monitor page by selecting **Monitor** > **Service Levels** from the left menu.

On the Monitor page, click the **WAN** button to view the WAN SLE Dashboard.



WAN Assurance Video Overview



Video:

Using the WAN SLE Dashboard

For a general introduction to the SLE dashboard, see ["Service-Level Expectations Overview" on page 24](#).

For help interpreting the WAN SLEs and classifiers, go to these topics:

- ["WAN Edge Health SLE" on page 55](#)
- ["WAN Link Health SLE" on page 57](#)

WAN Edge Health SLE

SUMMARY

Use the WAN Edge Health SLE to assess service levels for your WAN edge devices.

IN THIS SECTION

- [What Does the WAN Edge Health SLE Measure? | 55](#)
- [Classifiers for Poor WAN Edge Health | 55](#)
- [Example | 56](#)

WAN Edge Health is one of the WAN Service-Level Expectations (SLEs) that you can track on the WAN SLE dashboard in the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the WAN Edge Health SLE Measure?

Juniper Mist monitors the user minutes when the health or performance of the WAN edge device is not optimal. Suboptimal health lowers the device's ability to pass traffic, thus directly affecting any clients connected to the device.

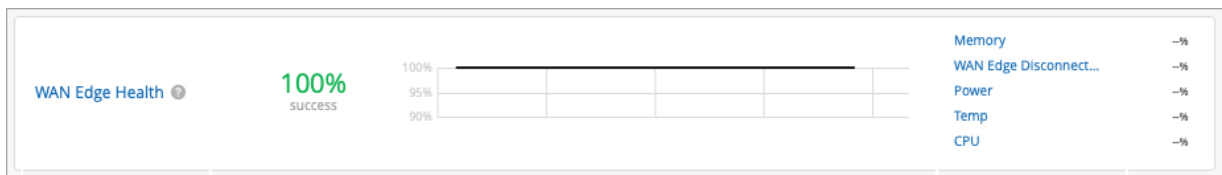
Juniper Mist analyzes various factors that affect WAN edge health and assigns a score. You can click the **Settings** button to set the success threshold.

Classifiers for Poor WAN Edge Health

When the WAN Edge Health threshold is not met, Juniper Mist classifies the issues as follows.

- **Memory**—Juniper Mist triggers this classifier when the WAN edge memory utilization is above 80 percent.
- **WAN Edge Disconnected**—Juniper Mist triggers this classifier when the WAN edge device disconnects from the Juniper Mist cloud.
- **Power**—Juniper Mist triggers this classifier when power consumption is above 90 percent of the available power.
- **Temperature**—Juniper Mist triggers this classifier when the operating temperature of the WAN edge device exceeds the prescribed threshold range, either going above the maximum limit or below the minimum requirement.
 - **CPU**—Juniper Mist triggers this sub-classifier when the CPU temperature exceeds the prescribed threshold range.
 - **Chassis**—Juniper Mist triggers this sub-classifier when the chassis temperature exceeds the prescribed threshold range.
- **CPU**—Juniper Mist triggers this classifier when the CPU utilization is above 90 percent. When the CPU utilization spikes on a Juniper WAN edge device, downstream devices can lose their connectivity. Therefore, clients fail to pass traffic.
 - **Data Plane**—Juniper Mist triggers this sub-classifier when the Data Plane CPU utilization is above 90 percent.
 - **Control Plane**—Juniper Mist triggers this sub-classifier when control plane CPU utilization is above 90 percent.

Example



- **Success Rate**—On the left, you see that the threshold was met 100 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see the list of classifiers for this SLE. There were no issues during the selected time period, so all these percentages are blank.

NOTE: If you click the **Values** button above the SLE blocks, you'll see numbers instead of percentages. Instead of success rate, you'll see the average severity level for the selected time period. In the Classifiers area, you see the total number of user minutes that were impacted by each classifier.

WAN Link Health SLE

SUMMARY

Use the WAN Link Health SLE to assess service levels for your WAN links.

IN THIS SECTION

- [What Does the WAN Link Health SLE Measure? | 57](#)
- [Classifiers for Poor WAN Link Health | 57](#)
- [Example | 58](#)

WAN Link Health is one of the WAN Service-Level Expectations (SLEs) that you can track on the WAN SLE dashboard in the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the WAN Link Health SLE Measure?

Juniper Mist monitors the user minutes when the WAN link health meets or fails to meet the SLE threshold. Poor WAN link health lowers the device's ability to pass traffic, thus directly affecting any clients using that link.

You can click the **Settings** button to set the success threshold.

Classifiers for Poor WAN Link Health

When the WAN Link threshold is not met, Juniper Mist classifies the issues as follows.

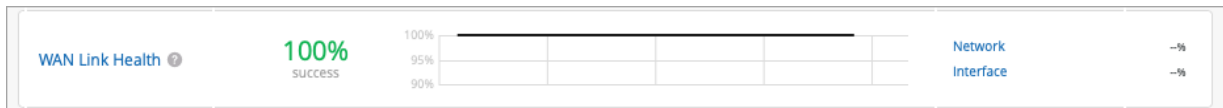
- **Network**—Network issues affected the WAN link.

The Network classifier has three sub-classifiers:

- **IPSec Tunnel Down**—One of the Overlay IPsec tunnels was down.

- **Latency**—WAN link traffic showed latency. Juniper Mist calculates latency by using the average value of round-trip time (RTT) for traffic over a period of time.
- **Jitter**—The WAN link experienced jitter. Juniper Mist calculates jitter by using the variation (standard deviation) of RTT within a period of 5 to 10 minutes for a particular WAN link. We compare the calculated value with the average deviation of RTT over a day or a week.
- **Interface**—Interface issues affected the WAN link. The Interface classifier has three sub-classifiers:
 - **Cable Issues**—Faulty cables affected the WAN link.
 - **Congestion**—Congestion affected the WAN link. The Congestion sub-classifier measures the number of output packet drops. When packets enter an interface, they go in a queue for buffering. When the buffer becomes full it starts to drop packets (Tx Drops).
 - **VPN**—VPN performance issue occurred.

Example



- **Success Rate**—On the left, you see that the threshold was met 100 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see the list of classifiers for this SLE. There were no issues during the selected time period, so all these percentages are blank.

NOTE: If you click the **Values** button above the SLE blocks, you'll see numbers instead of percentages. Instead of success rate, you'll see the average latency for the selected time period. In the Classifiers area, you'll see the total number of user minutes that were impacted by each classifier.

Location SLEs

SUMMARY

Use the Location service-level experience (SLE) dashboard to assess the service levels for user-impacting factors such as SDK connection issues, latency, dropped requests, access point health, and more.

IN THIS SECTION

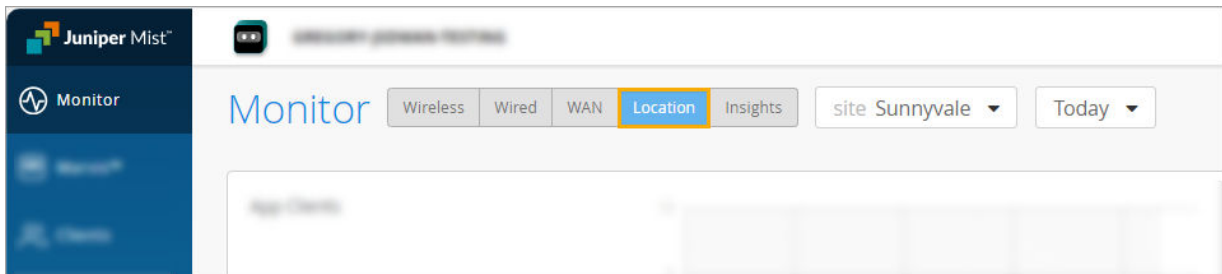
- [SDK Connect Time SLE | 60](#)
- [Location Latency SLE | 61](#)
- [Teleports SLE | 63](#)
- [Dropped Requests SLE | 64](#)
- [Location AP Health SLE | 67](#)

Finding the Location SLE Dashboard

The Location SLE Dashboard is part of the Monitor page, which appears automatically after you log in to the Juniper Mist™ portal.

From other pages of the portal, you can open the Monitor page by selecting **Monitor** > **Service Levels** from the left menu.

On the Monitor page, click the **Location** button to view the Location SLE Dashboard.



Using the Location SLE Dashboard

For a general introduction to the SLE dashboard, see ["Service-Level Expectations Overview" on page 24](#).

For help interpreting the location SLEs and classifiers, go to these topics:

- ["SDK Connect Time SLE" on page 60](#)
- ["Location Latency SLE" on page 61](#)
- ["Teleports SLE" on page 63](#)
- ["Dropped Requests SLE" on page 64](#)
- ["Location AP Health SLE" on page 67](#)

SDK Connect Time SLE

SUMMARY

Use the SDK Connect Time SLE to assess usage of your Juniper Mist™ SDK-enabled applications.

IN THIS SECTION

- [What Does the SDK Connect Time SLE Measure? | 60](#)
- [Classifiers for Low SDK Connect Time | 60](#)
- [Example | 61](#)

SDK Connect Time is one of the Location Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the SDK Connect Time SLE Measure?

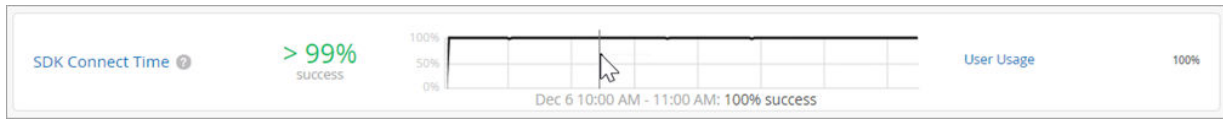
Juniper Mist measures the time when your SDK-enabled app clients are connected to location services at your site.

You can click the **Settings** button to set the number of seconds to use as the success threshold for this SLE.

Classifiers for Low SDK Connect Time

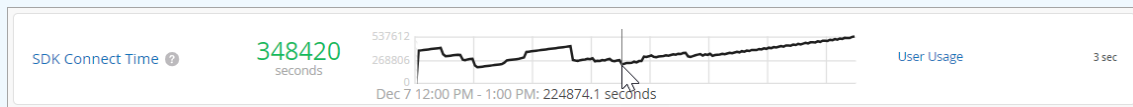
When the SDK Connect Time threshold is not met, Juniper Mist assigns one classifier: User Usage. Because there is only one classifier, it represents 100 percent of the app clients that did not meet the success threshold.

Example



- **Success Rate**—On the left, you see that the SDK Connect Time threshold was met by 99 percent of the app clients.
- **Timeline**—In the middle, you can drag your mouse to see the success rate at each point in time. This example shows 100 percent success at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, there is one classifier, which accounts for all app clients (100 percent) that failed to meet the usage threshold.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see the total number of user seconds for the selected time period. In the timeline, you see the number of user seconds for the selected time. On the right, you see that 3 user seconds failed to meet the SDK Connect Time threshold.



Location Latency SLE

SUMMARY

Use the Location Latency SLE to identify any user experiences that were affected by latency.

IN THIS SECTION

- [What Does the Location Latency SLE Measure? | 62](#)
- [Classifiers for Latency Issues | 62](#)
- [Example | 62](#)

The Latency SLE is one of the Location Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Location Latency SLE Measure?

Juniper Mist measures the latency of location responses and estimates to app clients.

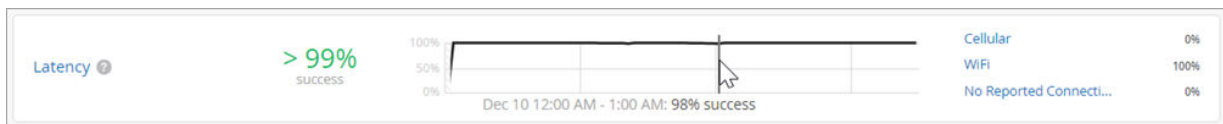
You can click the **Settings** button to set the number of milliseconds to use as the success threshold for this SLE.

Classifiers for Latency Issues

When the latency threshold is not met, Juniper Mist analyzes the data from the client devices and classifies the issues as follows.

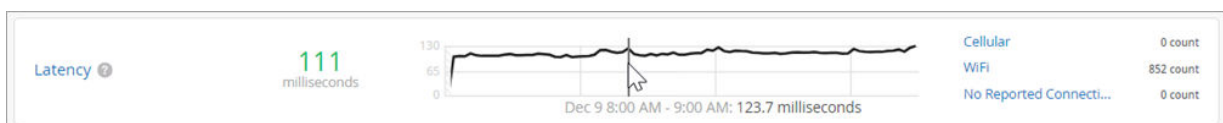
- **Cellular**—The app client is using a cellular data connection, which may attribute to the high latency.
- **WiFi**—The app client is using a wireless connection, which may attribute to the high latency.
- **No Reported Connection Type**—There is no information available about the connection type.

Example



- **Success Rate**—On the left, you see that the threshold was met 99 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. This example shows a 98 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the user experiences that failed to meet this SLE. All these issues (100 percent) were attributed to Wi-Fi. No issues (0 percent) were attributed to the other classifiers.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see the average latency in milliseconds for the selected time period. In the timeline, you see the latency in milliseconds at the selected time. On the right, you see the number of latency issues for each classifier.



Teleports SLE

SUMMARY

Use the Teleports SLE to identify issues with location accuracy.

IN THIS SECTION

- [What Does the Teleports SLE Measure? | 63](#)
- [Classifiers for Excessive Teleports | 63](#)
- [Example | 64](#)

Teleports is one of the Location Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Teleports SLE Measure?

Juniper Mist identifies instances when the app client's estimated location veers away (or "teleports") from the actual location.

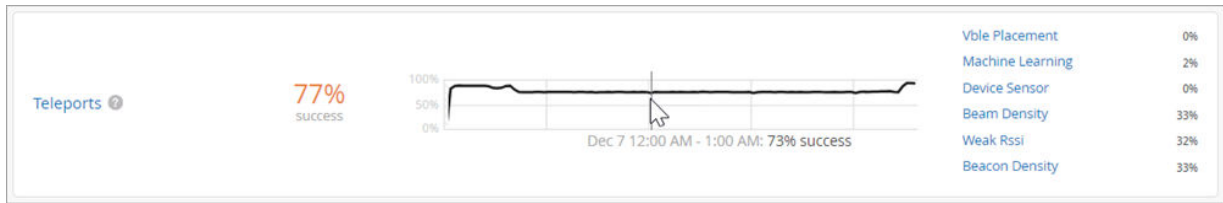
You can click the **Settings** button to set the number of meters to use as the success threshold for this SLE. The default is 3 meters, meaning that accuracy within 3 meters of the actual location is acceptable.

Classifiers for Excessive Teleports

When the Teleports SLE threshold is not met, Juniper Mist analyzes the data from the access points (APs) and the client devices and classifies the issues as follows.

- **Beacon Density**—The app client detected a low number of beacons from the access points (APs).
- **Beam Density**—The app client detected a low number of beams.
- **Machine Learning**—Changes in machine learning affected location accuracy.
- **vBLE Placement**—The placement of the APs affected location accuracy.
- **Device Sensor**—Sensor issues in the device affected location accuracy with respect to motion, acceleration, etc.
- **Weak RSSI**—The app client received a weak signal (low Received Signal Strength Indicator).

Example



- **Success Rate**—On the left, you see that the threshold was met 77 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to see the success rate at each point in time. The example shows a 73 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for teleports that exceeded the threshold. In this example, the issues were attributed to Beam Density (33 percent), Beacon Density (33 percent), Weak RSSI (32 percent), and Machine Learning (2 percent). No issues (0 percent) were attributed to the other classifiers (0 percent).

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see the average number of meters for teleports during the selected time period. In the middle, you see the number of meters for the selected moment. On the right, you see the number of excessive teleport issues per classifier.



Dropped Requests SLE

SUMMARY

Use the Dropped Requests SLE to identify issues with dropped location requests.

IN THIS SECTION

- [What Does the Dropped Requests SLE Measure? | 65](#)

●	Classifiers for Excessive Dropped Requests 65
●	Example 66

Dropped Requests is one of the Location Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Dropped Requests SLE Measure?

Juniper Mist monitors the instances when dropped location requests reduced location accuracy.

You can click the **Settings** button to set the number of pending requests to allow as the SLE success threshold. Pending requests are used as a likely indicator of requests that will be dropped in the future.

NOTE: The Pending Requests classifier uses the threshold that you set here. The other classifiers act on a pass/fail basing, counting any incidents that result in dropped requests.

Classifiers for Excessive Dropped Requests

When the number of dropped requests exceeds the threshold, Juniper Mist classifies the issues as follows.

- **Reconnects**—The app client briefly lost Internet connectivity and made attempts at reconnecting (successful or not).
- **Offline**—The app client went offline while using the app. The device might have had poor Wi-Fi reception, poor cellular reception, or connectivity problems. The user might have switched to airplane mode or turned off Wi-Fi.
- **Not Uniform Requests**—The app client sent some location requests at a faster or slower rate than expected, causing possible synchronization issues that affected location accuracy. (The app relies on uniform requests for proper location accuracy.)
- **Dropped by Network**—The app client's location requests were dropped while going through the network.
- **Client Request Timeout**—The app client's requests timed out.
- **Cellular**—The app client was using a cellular data connection.

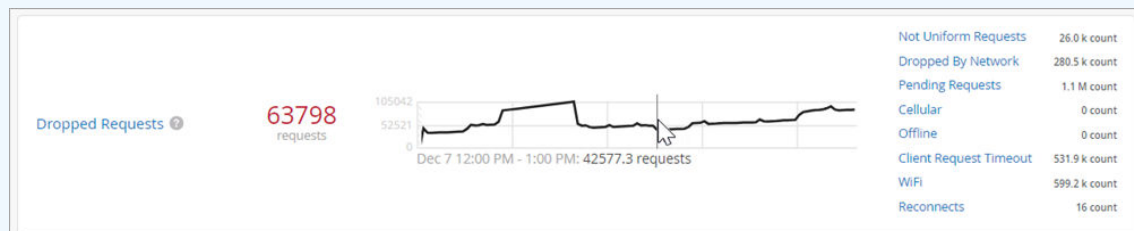
- **WiFi**—The app client was using a Wi-Fi connection.
- **Pending Requests**—The number of pending requests exceeded the SLE threshold.

Example



- **Success Rate**—On the left, you see that the Dropped Requests threshold was met only 6 percent of the time.
- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. This example shows an 11 percent success rate at the selected time. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the dropped requests. Many issues were attributed to Pending Requests (44 percent), WiFi (23 percent), and Client Request Timeout (21 percent). Other issues were attributed to Dropped by Network (11 percent), Not Uniform Requests (1 percent), and Reconnects (less than 1 percent). No issues (0 percent) were attributed to Cellular or Offline.

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see the number of dropped requests for the entire time period. The timeline shows the number of dropped requests at the selected time. On the right, you see the number of dropped requests for each classifier.



Location AP Health SLE

SUMMARY

Use the Location AP Health SLE to identify location issues caused by access points that rebooted or lost connectivity to the cloud.

IN THIS SECTION

- [What Does the Location AP Health SLE Measure? | 67](#)
- [Classifiers for Low AP Health | 67](#)
- [Example | 67](#)

Location AP Health is one of the Location Service-Level Expectations (SLEs) that you can track on the Monitor page of the Juniper Mist™ portal. Understand what's measured by this SLE and what issues can contribute to a low SLE.

What Does the Location AP Health SLE Measure?

Juniper Mist counts the incidents when APs rebooted or lost connectivity to the cloud.

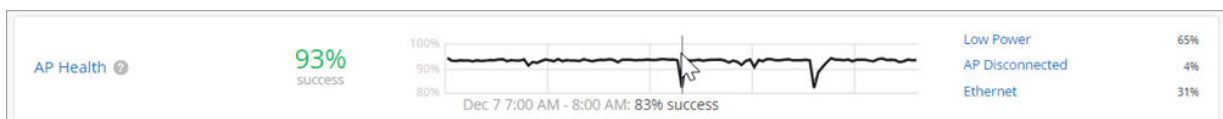
You can click the **Settings** button to set the number of incidents to allow as the success threshold for this SLE.

Classifiers for Low AP Health

When the number of incidents exceeds the threshold for this SLE, Juniper Mist classifies the issues as follows.

- **Low Power**—The AP had insufficient power.
- **AP Disconnected**—The AP had insufficient power.
- **Ethernet**—The AP lost Ethernet connectivity.

Example



- **Success Rate**—On the left, you see that the Location AP Health SLE threshold was met 93 percent of the time.

- **Timeline**—In the middle, you can drag your mouse to explore this SLE over time. In this example, the success rate was 83 percent at the selected moment. To adjust the scope of the timeline, use the timeline drop-down list at the top of the Monitor page. For example, set the timeline to Today, Yesterday, This Week, or a custom date range.
- **Classifiers**—On the right, you see a high-level root cause analysis for the incidents that failed to meet the threshold. Most (65 percent) were attributed to low power. The remaining issues were attributed to Ethernet (31 percent) and AP Disconnected (4 percent).

NOTE: You can click the **Values** button to show numbers instead of the success rate and classifier percentages. On the left, you see that there were 1371 incidents. In the middle, you see that there were 20.0 incidents at the selected moment. On the right, you see the number of user minutes that were attributed to each classifier.



4

CHAPTER

Alerts

[Alert Configuration](#) | 70

[Juniper Mist Alert Types](#) | 74

Alert Configuration

SUMMARY

Learn about alert configuration in the Juniper Mist dashboard.

IN THIS SECTION

- [Alerts Overview | 70](#)

Alerts Overview

IN THIS SECTION

- [Alert Configuration | 72](#)

In Juniper Mist, alerts present those events that don't fit neatly into the service-level experience (SLE) model. Whereas SLEs represent events that have already happened, alerts represent network and device issues that are ongoing. On the Monitor > Alerts dashboard, you can see three types of alerts: Infrastructure, Marvis, and Security.

Juniper Mist categorizes alerts that can potentially affect a large number of clients as infrastructure alerts. For example, an event during which a Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), or RADIUS server is unreachable can affect many clients. Similarly, if a power supply on a switch is in alarm state, a large number of clients and a large amount of traffic could be affected.

The Mist Predictive Analytics and Correlation Engine (PACE) raises Marvis alerts for the events that Marvis tracks. For example, if an access point (AP) regularly fails health checks, Marvis notices and tracks this event.

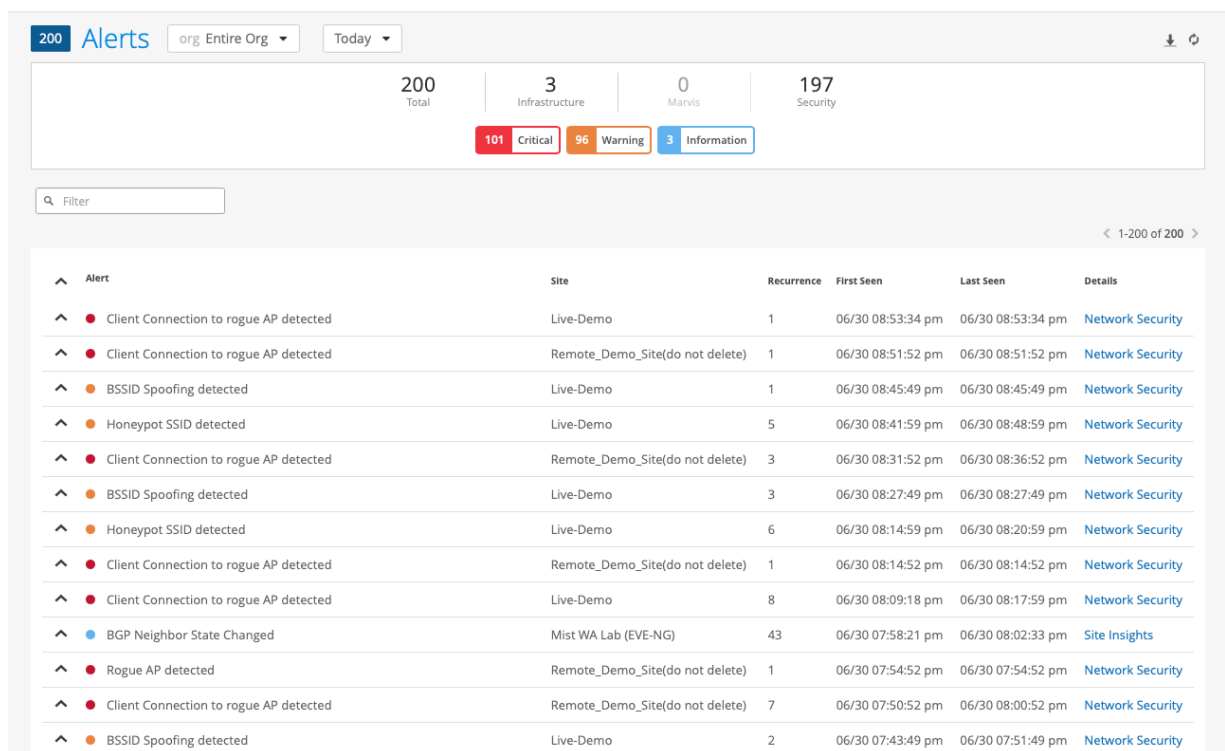
Security alerts are raised by repeated events that could dramatically affect network security. For example, if a rogue AP is detected, that represents a potential security problem and if a client connects to a rogue AP, that could be even worse.

We rank each category of alert by severity:

Table 1: Alert Severity

Severity	Indicator	Recommended Action
Critical	Red dot	Take immediate action.
Warning	Orange dot	Continue monitoring if the event continues.
Informational	Blue dot	No action is required.

In the image below, you can see an example of the Juniper Mist Alerts dashboard.



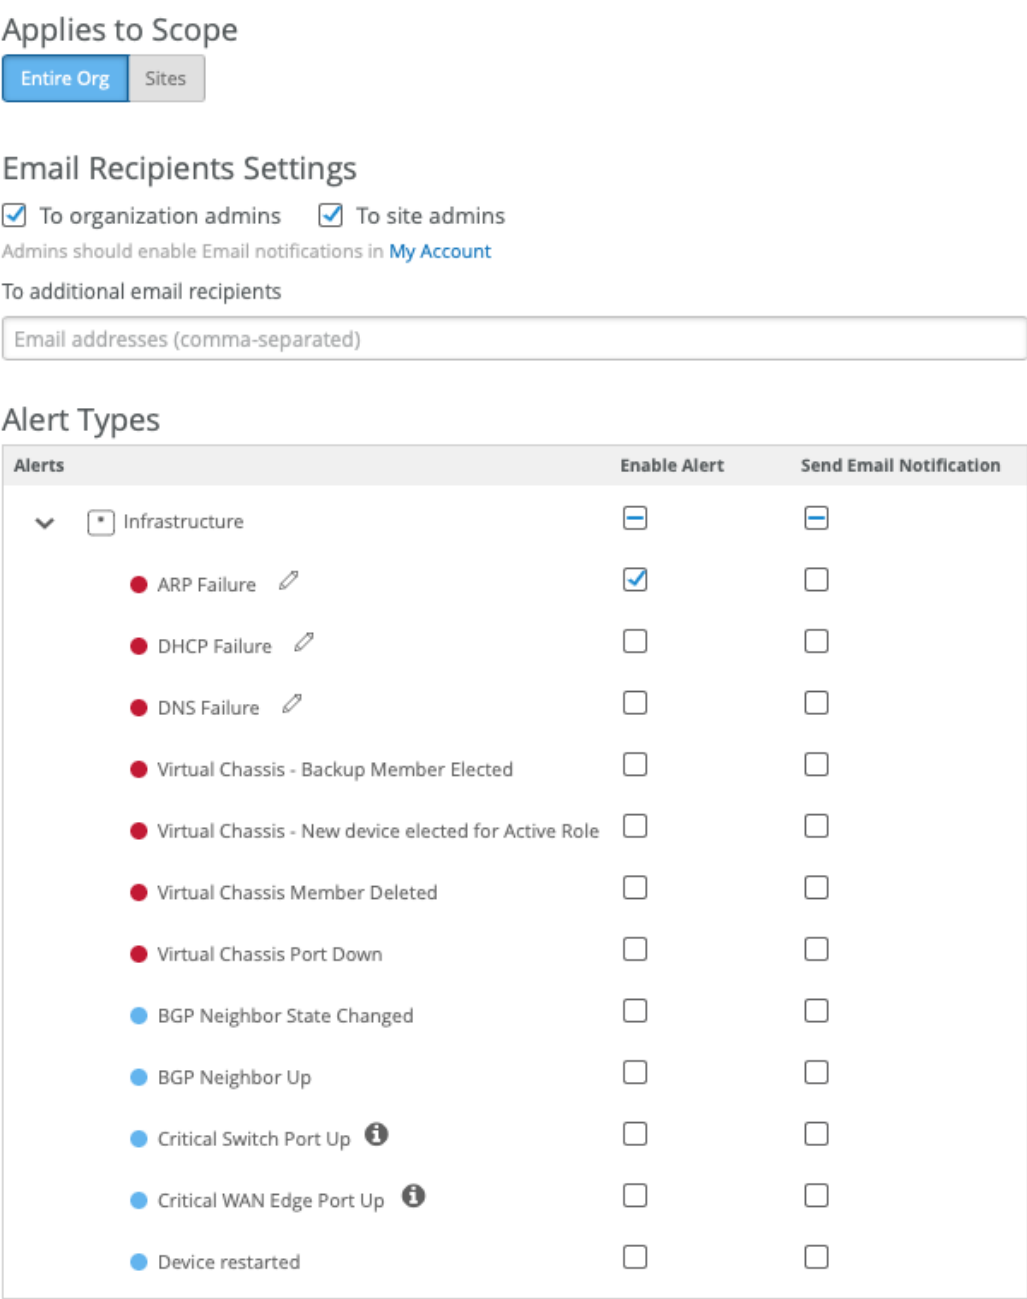
You can use the drop-down lists at the top of the Alerts dashboard to filter the dashboard display by:

- **Site**—You can choose to see alerts from your entire organization or from any site for which you have viewing rights. You can search within the menu for specific sites by name.
- **Date**—You can choose to see alerts from a specific date or time, or for a time range—from 60 minutes through 7 days.

Alert Configuration

When you click the **Alerts Configuration** button, you can view and configure Juniper Mist alerts. You can configure alerts on an organization-wide basis or on a site-by-site basis. You can enable or disable any alert by selecting or clearing the check box, respectively, next to the alert name. You can elect to have an e-mail notification sent to organization or site administrators, or to a comma-separated list of e-mail addresses. You can see an example of the [Figure 3 on page 73](#) page below.

Figure 3: Alert Configuration



on

page 73

In the [Figure 3 on page 73](#) image above, you can see:

- The alert configuration scope is set for the entire organization.
- E-mail notifications, if enabled, are configured to go to organization and site administrators.

- All the alert types belong to the Infrastructure category.
- None of the alert types are configured to send e-mail notifications.
- The critical *DNS Failure*, *DHCP Failure*, and *DNS Failure* alerts each have an edit (pencil) icon, which indicates that you can configure the threshold for each of these alerts.
- The informational *Critical Switch Port Up* and *Critical WAN Edge Port Up* alerts have an information (i in a circle) icon, which indicates that you must perform additional configuration to enable these alerts.

The two icons that indicate additional configuration options for alerts appear only for certain alerts in the infrastructure category. No alerts in other categories have additional configuration options.

Click the video below to watch the alert configuration process.



Video:

Juniper Mist Alert Types

IN THIS SECTION

- [Infrastructure Alerts | 74](#)
- [Marvis Alerts | 77](#)
- [Security Alerts | 79](#)

Infrastructure Alerts

In Juniper Mist, we present those events that don't fit neatly into the service-level experience (SLE) model as alerts. Whereas SLEs represent events that have already happened, alerts represent network and device issues that are ongoing. On the Monitor > Alerts dashboard, you can see three types of alerts: [Infrastructure on page 75](#) , [Marvis on page 77](#) , and [Security on page 79](#) .

Juniper Mist categorizes alerts that potentially affect a large number of clients as infrastructure alerts. For example, an event during which a Domain Name System (DNS), Dynamic Host Configuration

Protocol (DHCP), or RADIUS server is unreachable can affect many clients. Similarly, if a power supply on a switch is in alarm state, a large number of clients and a large amount of traffic could be affected.

The Mist Predictive Analytics and Correlation Engine (PACE) raises Marvis alerts for the events that Marvis tracks. For example, if an access point (AP) regularly fails health checks, Marvis notices and tracks this event.

Security alerts are raised by repeated events that could dramatically effect network security. For example, if a rogue AP is detected, that represents a potential security problem. If a client connects to a rogue AP, that could be even worse.

Table 2: Infrastructure Alerts by Severity

Severity	Alert Name	API Only
Critical	ARP Failure	
Critical	DHCP Failure	
Critical	DNS Failure	
Critical	Virtual Chassis - Backup Member Elected	
Critical	Virtual Chassis - New device elected for Active Role	
Critical	Virtual Chassis Member Deleted	
Critical	Virtual Chassis Port Down	
Informational	ARP Recovered	X
Informational	BGP Neighbor State Changed	
Informational	BGP Neighbor Up	
Informational	Critical Switch Port Up	
Informational	Critical WAN Edge Port Up	
Informational	Device reconnected	X
Informational	Device restarted	

Table 2: Infrastructure Alerts by Severity *(Continued)*

Severity	Alert Name	API Only
Informational	DHCP Recovered	X
Informational	DNS Recovered	X
Informational	HA Control Link Up	X
Informational	Switch reconnected	X
Informational	Switch restarted	
Informational	Virtual Chassis Member Added	
Informational	VPN Peer Up	
Informational	WAN Edge BGP Neighbor Up	
Informational	WAN Edge reconnected	x
Warning	BGP Neighbor Down	
Warning	Critical Switch Port Down	
Warning	Critical WAN Edge Port Down	
Warning	Device offline	
Warning	HA Control Link Down	
Warning	Loop detected (by AP)	
Warning	Switch Bad Optics	
Warning	Switch BPDU Error	
Warning	Switch DHCP Pool Exhausted	
Warning	Switch offline	

Table 2: Infrastructure Alerts by Severity *(Continued)*

Severity	Alert Name	API Only
Warning	Switch PEM Alarm	
Warning	Switch PoE Alarm	
Warning	Switch Power Supply Alarm	
Warning	Switch Storage Partition Alarm	
Warning	Tunnel down	
Warning	VPN Peer Down	
Warning	WAN Edge BGP Neighbor Down	
Warning	WAN Edge DHCP Pool Exhausted	
Warning	WAN Edge offline	x
Warning	WAN Edge Source NAT Pool Threshold Exceeded	

Marvis Alerts

Marvis alerts are tied into the **Marvis Action Dashboard**. These alerts trigger whenever the corresponding Marvis Action is detected in your organization. If an AP regularly fails health checks, Marvis notices and tracks it.

The table below provides a listing of Marvis alerts, sorted by severity.

Table 3: Marvis Alerts by Severity

Severity	Applies To	Alert Name
Critical	AP	AP health check failed
Critical	AP	AP insufficient capacity

Critical	AP	AP insufficient coverage
Critical	AP	Bad cable
Critical	AP	Non-compliant
Critical	AP	Offline (Marvis)
Critical	connectivity	ARP failure (Marvis)
Critical	connectivity	Authentication failure (Marvis)
Critical	connectivity	DHCP failure (Marvis)
Critical	connectivity	DNS failure (Marvis)
Critical	WAN edge	Bad cable
Critical	WAN edge	Bad WAN Uplink
Critical	WAN edge	Negotiation mismatch
Critical	WAN edge	VPN Path Down
Critical	switch	Bad cable
Critical	switch	Missing VLAN
Critical	switch	Negotiation mismatch
Critical	switch	Port Stuck
Critical	switch	Switch STP Loop
Warning	switch	Port flap

Security Alerts

Security alerts warn you of activities or events on the network that can cost you in terms of lost data, unauthorized access to the network, or traffic that matches known security threats. Juniper Mist lists all security alerts except those that relate to intrusion detection and prevention (IDP) or URL filtering on the Monitor > Alerts page. You can find IDP and URL filtering events and their severity on the **Site > WAN Edge > Secure WAN Edge IDP/URL Events** page.

Table 4: Security Alerts by Severity

Severity	Alert Name
Critical	Client Connection to rogue AP detected
Critical	Rogue AP detected
Informational	Air Magnet Scan detected
Informational	EAP Handshake Flood detected
Warning	Active Watched Station detected
Warning	Adhoc Network detected
Warning	BSSID Spoofing detected
Warning	Disassociation Attack detected
Warning	EAP Dictionary Attack detected
Warning	EAP Failure Injection detected
Warning	EAP Spoofed Success detected
Warning	EAPOL-Logoff Attack detected
Warning	ESSID Jack detected
Warning	Excessive Clients detected
Warning	Excessive EAPOL-Start detected

Warning	Fake AP Flooding detected
Warning	Honeypot SSID detected
Warning	IDP attack detected
Warning	Monkey Jack detected
Warning	Out of Sequence detected
Warning	Repeated Client Authentication Failures
Warning	Replay Injection detected - KRACK Attack
Warning	Security Policy Violation
Warning	SSID Injection detected
Warning	TKIP ICV Attack
Warning	URL blocked
Warning	Vendor IE Missing
Warning	Zero SSID Association Request detected

RELATED DOCUMENTATION

No Link Title

[Alert Configuration](#) | 70