Juniper Networks | Engineering Simplicity

MOMENTUM

Juniper Mist Network Monitoring

Published
2023-07-31

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

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
CHAPTER

Overview

Network Monitoring with Juniper Mist

IN THIS SECTION

Read this section to learn about network and device monitoring capabilities in the Juniper Mist portal.

What Is Network Monitoring?

Network monitoring encompasses the tasks that you must perform on the Juniper Mist portal to check network health. To monitor your network, you must:

- Track network device health, performance, and status.
- Know the amount and types of traffic passing through the network.
- Know the clients attached to the network along with client health and traffic information.

With this kind of monitoring capability, you can spot trends in the network, identify devices that consistently have issues, and even find the bad cables connected to your wired switches.

When you use the Juniper Mist portal to monitor your network, you gain insight into what's happening often before it becomes an issue. You can see the network from multiple perspectives: wireless, wired, and WAN. You can configure alert levels and thresholds to suit your needs. Additionally, you can correct potential issues with the tools that Juniper Mist provides for managing AP, switch, WAN edge, and Juniper Mist Edge congfiguration.

Mist Insights

Juniper Mist™ monitors connection, status, traffic throughput, latency, event, and health information for:

- Juniper Mist Access Points (APs)
- Wireless client devices
- Wired switches
- Wired client devices
- Network servers
- Network applications
- WAN edge devices

Juniper Mist then compiles and correlates this information into dashboard views called Insights. The Juniper Mist portal provides dashboard views for sites, access points, and clients. Mist Insights presents network and device event information on the Monitor > Service Levels dashboard when you select the **Insights** tab. Juniper Mist categorizes the events as good, bad, or neutral and displays the individual events with green, red, or orange highlights, respectively.

Wireless, Wired, WAN, and Location insights each have their own specific dasboards. You can access these specific dashboards when you select the appropriate tab on the Monitor > Service Levels page. An example of the Insights dashboard for Sites is shown below.

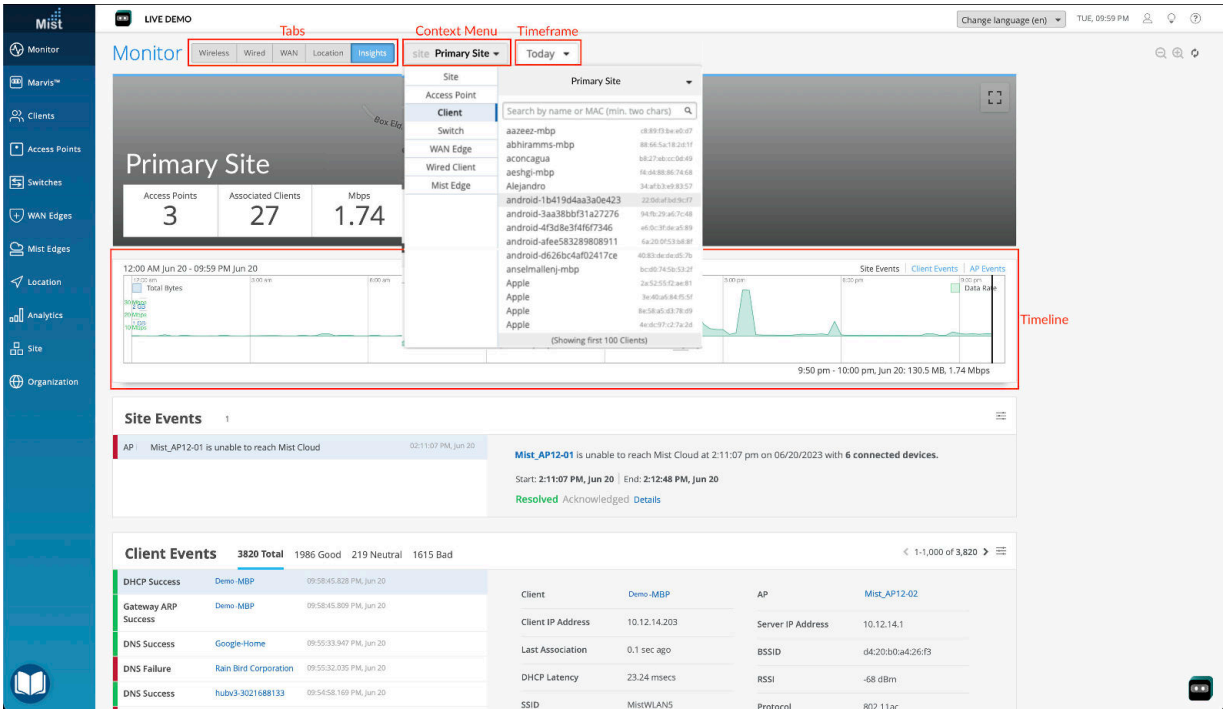

See "Site Insights" on page 7 for more information.

Mist Alerts

Juniper Mist uses alerts to notify you about network and device events that occur in your organization or individual sites. The Monitor > Alerts dashboard helps you maintain full and constant visibility across your organization and sites. On this page, you can view existing alerts and configure the display and notification properties of alerts throughout your organization. While the Mist Insights dashboard provides a lot of data about events that have occurred in each section, the Monitor > Alerts dashboard provides precise information about individual device or ongoing network service events. While viewing a specific alert, you can acknowledge that you have seen the alert. You can also configure a filter on the Alerts dashboard to show only unacknowledged alerts. An example of the Alerts dashboard is shown below:

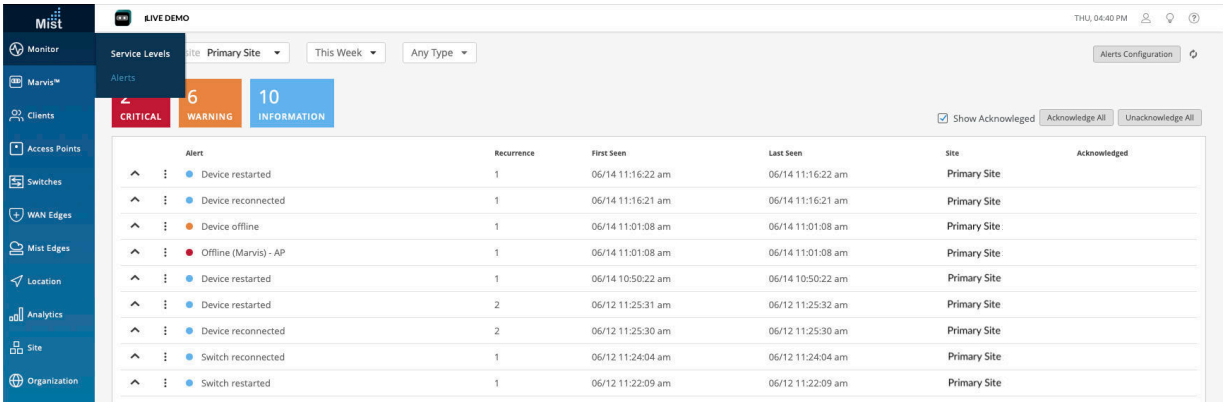

See "Alert Configuration" on page 39 for more information.

Marvis Actions

In addition to insights and alerts, Juniper Mist uses the Marvis AI to both act on your behalf and notify you of actions you can take. Marvis can identify the root cause of certain issues and recommend corrective actions to you. On the Marvis page, you can see the actions that Marvis has identified or performed as a result of the activity on the network. You can also query Marvis to understand many network and device behaviors. An example of the Marvis Actions page is shown below.

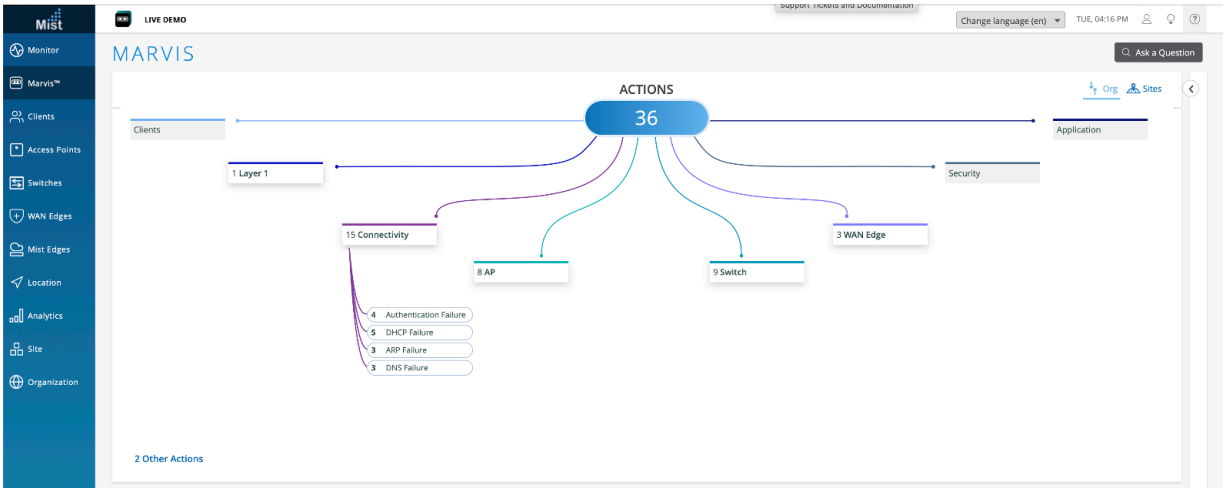

See Marvis Actions for more details.

How User Roles Affect Network Monitoring

You can configure user roles to provide read-write or read-only access to the monitoring features in the Mist platform. When you assign an admistrator a read-only role, they can only view alerts and insights. Read-only administrators cannot acknowledge or make any configuration changes to alerts or insights. See *Add Accounts*and *Portal User Roles* in the Mist Management Guide for details about how you configure administrators and roles.

.You must enable your account to receive email notifications from the Monitor > Alerts dashboard.

1

PART

Insights

Mist Insights

IN THIS SECTION

Site Insights

In the Juniper Mist portal, the **Insights** tab on the **Monitor > Insights Service Levels** page provides high-level awareness of what is happening in your network. The Mist Predictive Analytics and Correlation Engine (PACE) drives the Juniper Mist insights and service level experiences (SLEs). The Mist PACE collects:

- Telemetry data from Juniper wired switches, WAN edge devices, and Juniper Mist Edge devices
- Time to connect data from wireless clients
- Coverage, roaming, and throughput data from Juniper Mist access points (APs)
- Throughput data for network applications

- Dwell time and other location data from Bluetooth Low Energy (BLE) tags

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

You can change the context of the displayed insights from an entire site to individual access points (APs) or clients by selecting the option from the context menu. In addition, you can view insights for the last 60 minutes through the last 7 days.

Figure 1: Monitor Insights Dashboard with Options

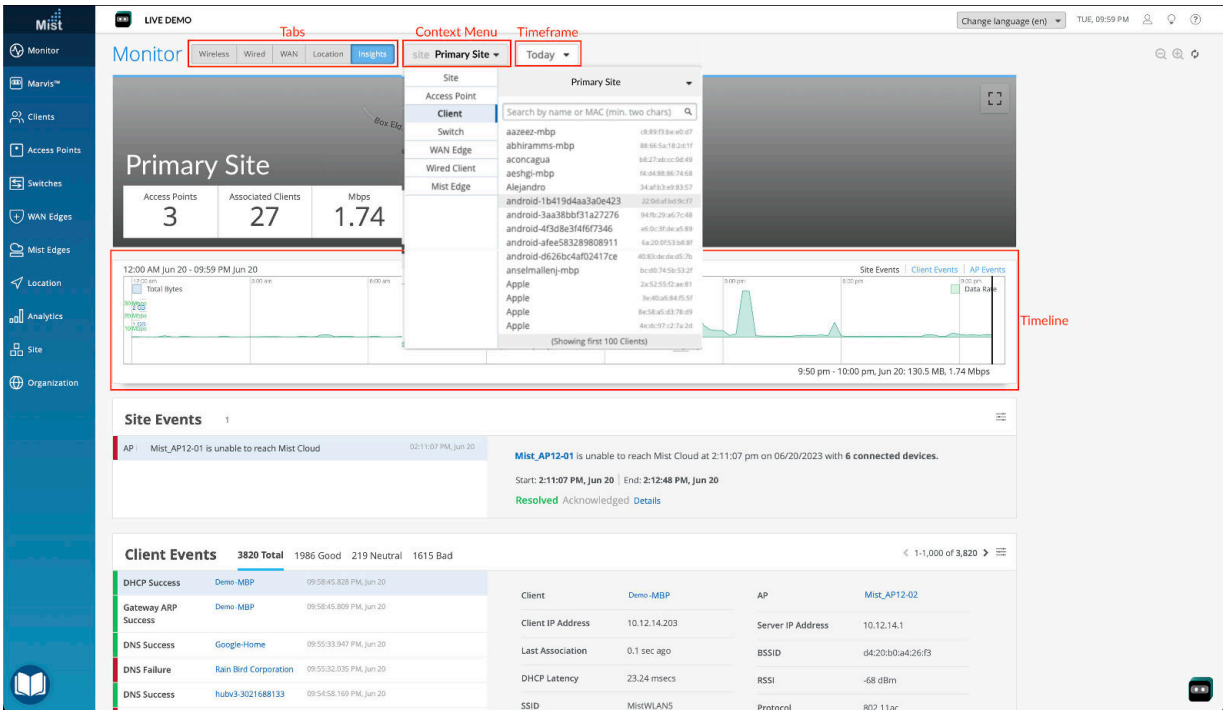

When you navigate to **Monitor > Service Levels**, you see the Insights dashboard. The first time you navigate to this page, the Insights tab is selected by default. Also by default, the context menu is set to the primary site within your organization (site context). A search field is available in the context menu regardless of which context you select. When set to the Site context, the Insights dashboard displays the following sections below the timeline:

- **Site Events**—Identify events that affect a large number of devices or clients. Site events include such things as DNS or DHCP server reachability or AP reboot events. Currently, site events apply only to wireless networks.

- **Client Events**—Learn about events that are related to or reported by individual client devices. The reporting clients can be wired or wireless. An example of a client event is when a DNS request from a client fails.
- **AP Events**—Determine events that are related to or reported by a AP. An example of an AP event is when the configuration of an AP is changed.
- **Applications**—Discover the network applications that are running on your network. The application insights that this section provides include the percentage of traffic using a given application and the total bytes of data on the network for each application.
- **Network Servers**—Get information for three network server types: RADIUS, DHCP, and DNS.
- **Pre-connection and Post-connection**—Know how long it takes for a wireless client to connect to the network. Use these insights to determine the connected client count and amount of data sent and received.
- **Current Site Properties**—Know more about the site location properties, including information about the current count of devices and clients.
- **Current WLANs**—View the count and names of all wireless LANs (WLANs) defined on your site along with information for each wireless LAN.
- **Access Points**—Get the current status information about all access points configured at the site.
- **Clients** —Discover information about the count, name, and types of wireless clients associated with the current site.
- **Wired Switches**—know about any wired switches associated with the site. The information includes the switch name, the IP address, and a count of APs connected to the switch.

NOTE: The sections displayed on the Insights dashboard depend on the subscriptions you have licensed. The list above represents the sections available for a fully licensed organization. Throughout this topic, we describe the features and screens available to a fully licensed organization. If you have purchased only a Wireless Assurance license, then Juniper Mist displays only the Wireless and Insights tabs on the Monitor > Service Levels dashboard.

Context Menu

The context menu lets you select the context for the dashboard display. The menu provides eight separate contexts to choose from:

- Site
- Access Point
- Client
- Switch
- WAN Edge
- Wired Client
- Mist Edge
- Cellular Edge

Site is the broadest context available. You can select any site configured within your organization. The default is your primary site or the first site created. You can search for a site by its name if you have a large number of sites.

The remaining context options enable you to display insights for specific devices or clients. By default the menu displays the APs, clients, switches, etc. associated with the primary site. If you select a different site, the menu displays the APs, clients, switches, etc. associated only with that site. While the menu is displayed, you can change the site context by selecting a specific site name or 'Any Site.'

You can search within the menu for devices or clients by name or MAC address.

Map Display

The Insights dashboard displays a map directly below the tabs. The map displayed depends on the context selected in the context menu. If you select Site, the map shows the geographic location of the site. If you select Access Point, the map displays the location of the AP within the site map. When you select other contexts, the map changes to display that item's location relative to the selected site or geography.

NOTE: If you have not set up a map of your site location, the map will display as a blank, grey block with the site name overlayed in white text.

Insights Timeline (Time Range)

The time range graph that appears directly below the map displays only today's events by default. You can pull down the **Today** menu to change the time range display. As mentioned previously, you can zoom the display in as close as the last 60 minutes or zoom the display out as far as the last 7 days.

You can change the time range to display Site Events, Client Events, or AP events by clicking the appropriate blue text in the time range.

If you purchase a Premium Analytics subscription, you can access up to 3 years' worth of wireless network insights and other data. You can see the data for a longer time range only through the Premium Analytics interface, not on the Monitor > Service Levels > Insights dashboard.

Site Events

The Site Events block displays 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 that Mist performed as a result of this event such as sending email or SMS messages.
- 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 Mist 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

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. We put a device at the center of a circle. The signal strength is strongest (best) when you are close to the device. As you move toward the outside of the circle, the signal strength diminishes. The heat map displays APs 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 represnts 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 strenght 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.

RELATED DOCUMENTATION

Network Monitoring with Juniper Mist | 2

Alert Configuration | 39

2

PART

Service Level Experiences

Service Level Experiences (SLEs) Overview

IN THIS SECTION

Overview

The Mist Predictive Analytics and Correlation Engine (PACE) analyzes, correlates, and classifies event and performance data captured from the network and its devices. PACE uses the resulting data to provide you with a view of user experience called Service Level Experience (SLE). Detailed SLE metrics are presented in dashboard views for four network service areas:

(Click any link in the list below to go directly to the details of the SLEs for that network service and their contributing classifiers.)

- "Wireless" on page 20
- "Wired" on page 28
- "WAN" on page 33

You access the dashboards by navigating to **Monitor** → **Service Levels** and selecting one of the network service area tabs. For example, an altered view of the Figure 3 on page 18 is shown below.

Figure 3: The Wireless SLE Dashboard

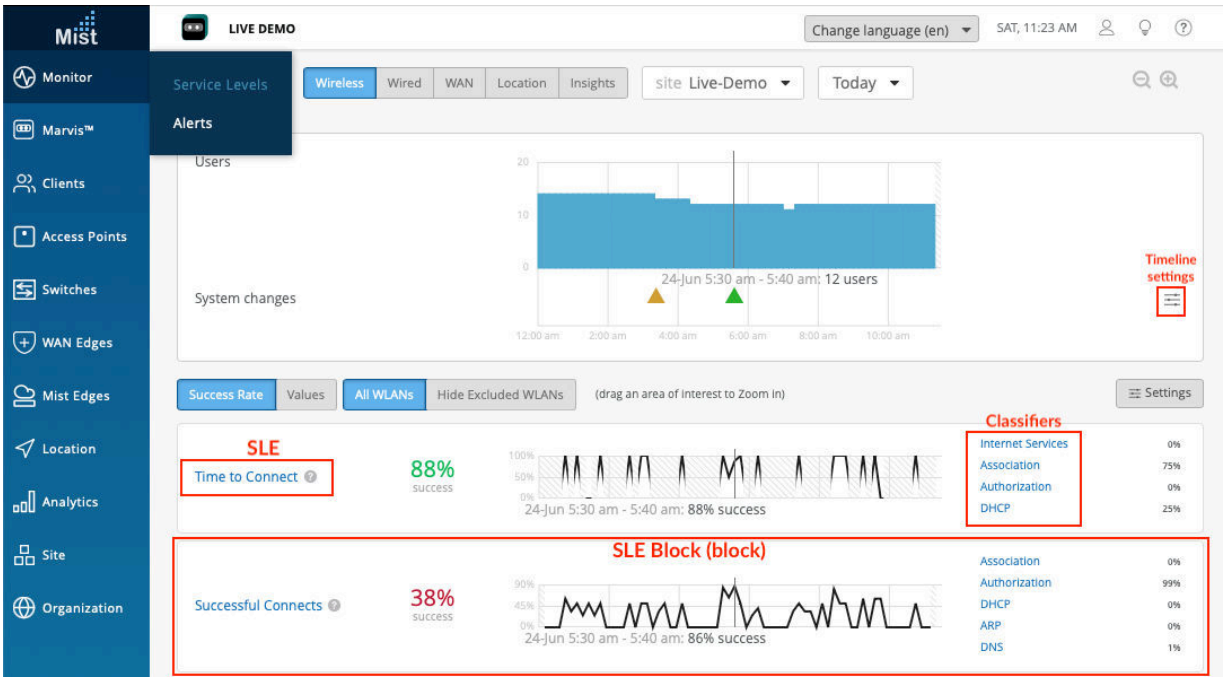

The Figure 3 on page 18 above only shows two of the available Wireless SLEs: *Time to Connect* and *Successful Connects*. In this example, the *Successful Connects* metric shows that only 38% of the connection attempts were successful. This means that 62% were unsuccessful. The list of classifiers shows you that 99% of the time, the connection problem was related to *Authorization*.

Later in this topic, we dive into the details of the SLEs and Classifiers for each network service area.

Referring back to Figure 3 on page 18, you can see the timeline with the green and orange triangles below. The triangles represent system changes. Click the settings button inside the timeline to understand the meaning of the triangles and to filter the kinds of events that make System Changes appear below the timeline.

Mist displays SLE metrics as Success Rate (percentage) by default. You can change this to show Values instead.

Click the **Settings** button to filter which SLEs display on the dashboard and in what order. The Figure 4 on page 19 pop-up shows the SLEs available for the Wireless service area. We have selected the Throughput SLE. Drag the orange line left or right to adjust the throughput target for the wireless network. Not all SLEs have adjustable targets or thresholds.

Figure 4: Wireless SLE Settings

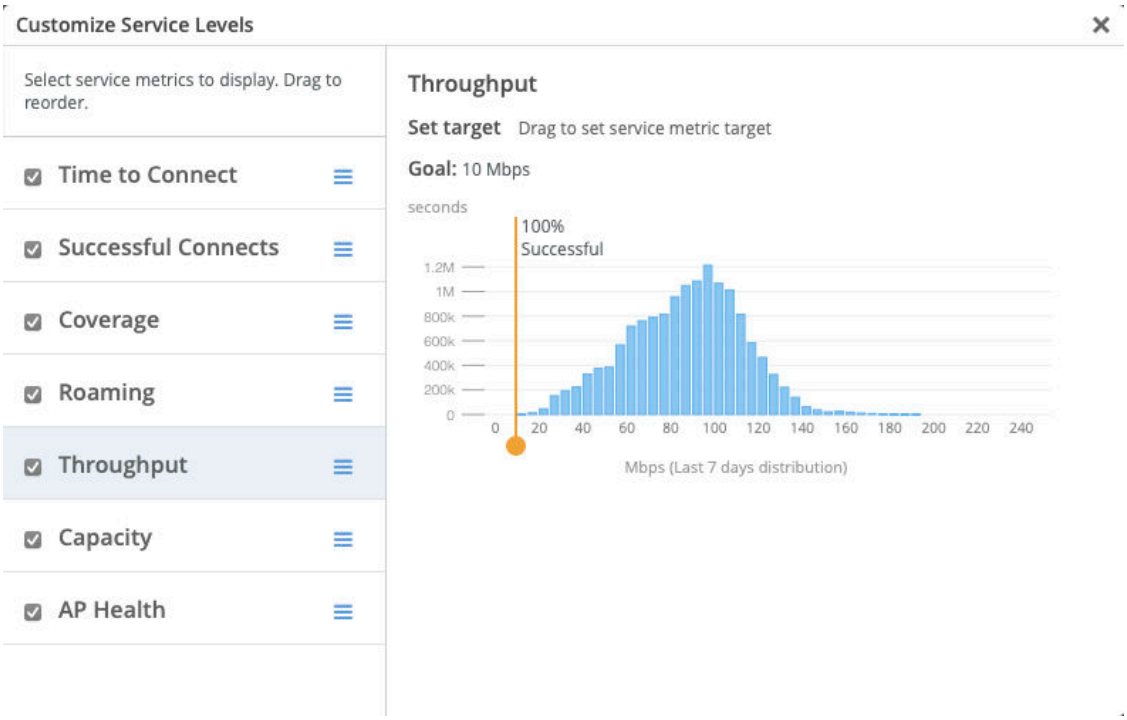

When you hover over the question mark or the blue text in any of the SLE blocks, Mist displays a brief definition of the SLE or Classifier. When you click the blue text in any of the SLE blocks, Mist takes you to the Root Cause analysis page for that SLE. There you can see which classifiers PACE used to determine the rating percentage in the SLE block. Some classifiers have sub-classifiers that also contribute to the rating. You see these sub-classifiers when you click on a classifer name on the Root Cause analysis page. Below, you can see an example of the Root Cause analysis page for the Wireless/ Time to Connect SLE.

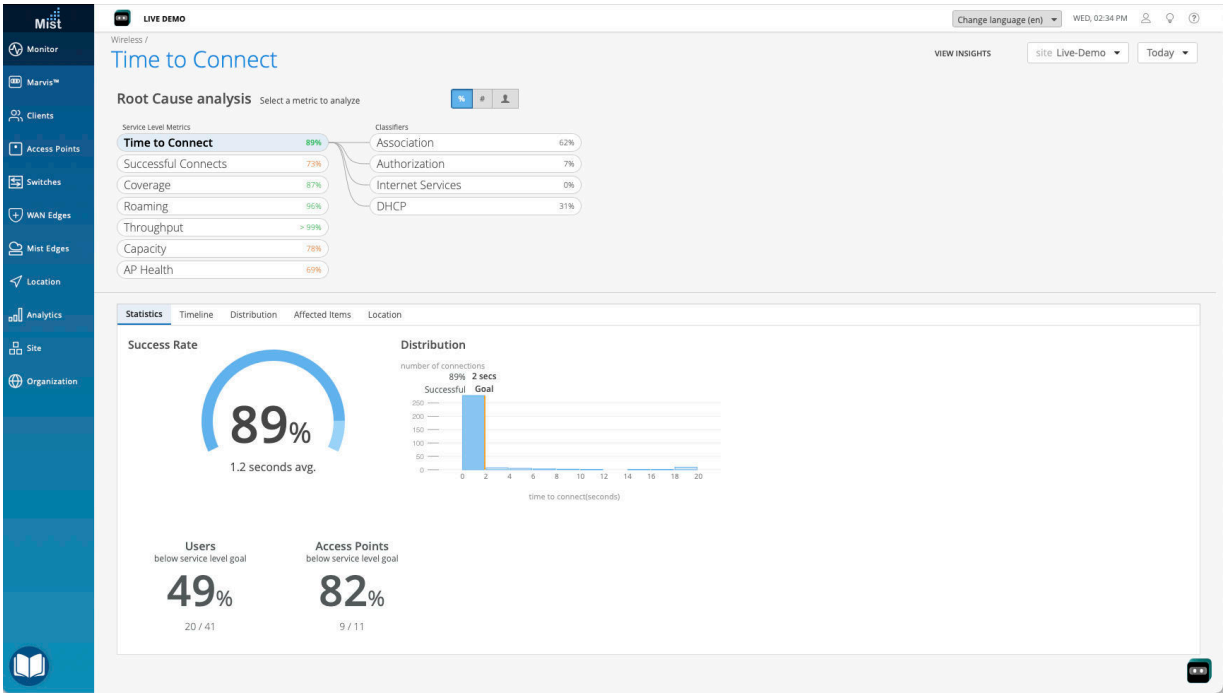

The remainder of this topic provides details about the SLE categories and the classifiers and sub-classifiers in each category.

Wireless SLEs

IN THIS SECTION

Overview

The wireless SLE dashboards display the percentage of time that the SLE metrics met the specified service level expectation goal within a specific time range. These metrics are categorized into classifiers and sub-classifiers, which provide additional details to identify the specific causes of failure. With this information, you can easily identify and address the issues affecting the end-user experience.

Our wireless SLEs are divided into 7 Metrics: **Time to Connect**, **Successful Connects**, **Coverage**, **Roaming**, **Throughput**, **Capacity** and **AP Health**. Click into these metrics to find detailed Classifiers and Sub-Classifiers which show the exact circumstances of failures experienced on your wireless network.

The table, Table 1 on page 21, below shows the names of the Wireless SLE metrics along with the classifiers and sub-classifiers for that metric. Click on the metric name to see a definition of the metric and its classifiers.

Table 1: Wireless SLEs, Classifers and Sub-Classifiers

Metric Name	Classifiers	Sub-Classifiers
"Time to Connect" on page 23	Internet Services	n/a
	Authorization	n/a
	Association	n/a
	DHCP	Unresponsive
		Stuck
		Nack
" Successful Connects " on page 24	Association	n/a
	Authorization	n/a
	DHCP	Discover Unresponsive
		Renew Unresponsive
		Incomplete
		Nack

Table 1: Wireless SLEs, Classifers and Sub-Classifiers *(Continued)*

Metric Name	Classifiers	Sub-Classifiers
	ARP	n/a
	DNS	n/a
" Coverage" on page 25	Asymmetry Downlink	n/a
	Weak Signal	n/a
	Asymmetry Uplink	n/a
" Roaming" on page 25	Latency	Slow 11r Roam
		Slow OKC Roam
		Slow Standard Roam
	Stability	Failed to Fast Roam
	Signal Quality	Interband Roam
		Suboptimal Roam
		Sticky Client
"Throughput" on page 26	Coverage	n/a
	Network Issues	n/a
	Device Capacity	n/a
	Capacity	Non WiFi Interference
		WiFi Interference
		High Bandwidth Utilization
		Excessive Client Load

Table 1: Wireless SLEs, Classifers and Sub-Classifiers *(Continued)*

Metric Name	Classifiers	Sub-Classifiers
"Capacity " on page 27	Non WiFi Interference	n/a
	Client Usage	n/a
	Client Count	n/a
	WiFi interference	n/a
"AP Health" on page 28	Low Power	n/a
	AP Disconnected	Switch Down
		Site Down
		AP Reboot
		AP Unreachable
	Ethernet	Speed Mismatch
		Ethernet Errors

Time to Connect

This SLE metric tracks the number of connections that took longer than the specified threshold to connect to the Internet. The time to connect to the Internet is calculated as the time from the start of the association packet from the mobile client to the point where the client can successfully move data.

time_to_connect = tconnected – tfirst-assoc

The classifiers for the time to connect metric are triggered if the time to connect exceeds the specified threshold. If the client does not connect to the Internet, this metric does not count the connection towards the connect time metric. The Connect Time is tracked by a separate service level metric. The current implementation has the classifiers divide up the time_to_connect into various buckets.

time_to_connect = sum(all tclassifier)

Time to Connect Classifiers

- **Internet Services** - This classifier is triggered if the users time to access external networks is more than 2 sigma from the moving average for this site.
- **Authorization** - This classifier is assigned if a user's time to go past the "authentication" state is more than 2 sigma from the average authentication latency, for this site.
- **Association** - This classifier is assigned if a users's time to go past the "association" state is more than 2 sigma from the average association latency, for this site
- **DHCP** - This classifier is assigned if a user's DHCP time is more than 2 sigma from the average DCHP time of fully completed successful connections for this site.
    - Stuck
    - Nack
    - Unresponsive

Back to SLE list on page 21

Successful Connects

This SLE metric tracks the percentage of successful Authorization, Association, DHCP, ARP, and DNS attempts during the initial connection by a client to the wireless network when a client roams from one AP to the next on an on-going basis.

Successful Connects Classifiers

- **Association** – This classifier is assigned if the connection fails during the Association process.
- **Authorization** – This classifier is assigned if the connection fails during the Authorization process.
- **DHCP** – This classifier is assigned if the connection fails during the DHCP process. This classifier contains four sub-classifiers:
    - Renew Unresponsive
    - Nack
    - Incomplete
    - Discover Unresponsive
- **ARP**– This classifier is assigned, if during the connection, ARP for the default gateway fails, or if ARP gateway failures are experienced after the initial connection or roam.

- **DNS** – This classifier is assigned if DNS failures are experienced during or after the connection process.

Back to SLE list on page 21

Coverage

This SLE metric tracks the number of user minutes that a client's RSSI, as measured by the access point, is below the threshold configurable by IT. This metric accounts for client activity. If the client is not active, the classifiers are not fired. The Asymmetry classifiers display bad coverage between the client and the AP. This field is usually displayed in minutes (Number of minutes of bad coverage). Asymmetry indicates there is a power level mismatch between the client & AP. Since APs are capable of higher power levels than a client, if power levels are set too high, then there is a chance that the low powered client is not heard by the AP, resulting in an asymmetry uplink issue. Asymmetry measurements are divided into 2 Classifiers: Asymmetry Uplink and Asymmetry Downlink.

Coverage Classifiers

- **Asymmetry Uplink** – This classifier tracks the number of user minutes that a client experiences bad coverage that can be attributed to asymmetric uplink transmit powers between the AP and client device. Meaning, the AP hears a weak signal from the clients for a number of minutes. There are various reasons causing this, like clients being too far from the AP. The traffic going from the client to the AP, and then to Internet is called uplink traffic.
- **Asymmetry Downlink** – This classifier tracks the number of user minutes that a client experiences bad coverage that can be attributed to asymmetric downlink transmit powers between the AP and client device. Meaning, clients are hearing a weak signal from the AP. The traffic going from the AP to the client is called downlink traffic.
- **Weak Signal** – This classifier tracks all other user minutes below the RSSI threshold.

Back to SLE list on page 21

Roaming

This SLE metric tracks the percentage of successful roams between 2 access points for clients that are within the prescribed thresholds. The user defines the threshold as a target time it takes for a client to roam. Fast roaming as defined by 802.11r and OKC are for clients using RADIUS based authentication.

Additionally, the Roaming SLE success threshold has changed from time based to severity based. Severity is the score from 1 to 5, where 1 is excellent roaming in the site and 5 is poor roaming in the site. The default is a severity score of 2.

Roaming Classifiers

- **Latency** – This classifier tracks the delta time between clients roaming across APs.
  - **Slow 11r Roams** – Roam time exceeds 400 ms.
  - **Slow Standard Roams** – Roam time exceeds two seconds.
  - **Slow OKC Roams** – Roam time exceeds 400 ms.
- **Stability** – This classifier tracks the consistency of AP choice and 11r usage during client roams. This classifier is assigned if a fast roam capable user on a fast roam enabled SSID experienced a slow roam that took more than 2 seconds. The stability classifier contains one sub-classifier: **Failed to fast Roam**.
- **Signal Quality** – This classifier tracks the RSSI of clients during a roam event.
  - **Suboptimal Roam** – Suboptimal Roam tracks when clients roam to an AP with more than a 6 dB decrease in RSSI from the previous AP, and if the new association RSSI is worse than the configured coverage SLE threshold (default -72 dBm).
  - **Sticky Client** – Sticky Client tracks when a client remains connected to an AP even when there are more roaming options available to improve the RSSI by more than 6 dB.

Back to SLE list on page 21

Throughput

This SLE metric tracks the amount of time that a client's estimated throughput is below the target threshold that you configure on the Figure 4 on page 19 page.

A client's estimated throughput is defined as the probabilistic throughput given the client's current wireless conditions. The estimator considers effects such as AP bandwidth, load, interference events, the type of wireless device, signal strength, and wired bandwidth. The estimated throughput is calculated on a per-client basis for the entire site.

Throughput Classifiers

PACE uses four classifiers for low throughput. These four are the likely causes for potential low throughput.

- **Device Capability** – This metric tracks the user minutes that client's throughput is below the configured threshold, primarily due to the capacity of device.

- **Capacity** – This classifier tracks the user minutes that the client's throughput is below the configured threshold, due to the load on the associated access point and Wi-Fi/ Non-Wi-Fi interference on the channel. The capacity classifier has four sub-classifiers:
  - High Bandwidth Utilization
  - Non WiFi Interference
  - Excessive Client Load
  - WiFi Interference

  In these sub-classifiers, you can examine Users and APs below the service level goal, the Timeline of failures and system changes, the distribution of failures, and affected items relating to the sub-classifier.
- **Coverage** – This metric tracks the user minutes that client's throughput is below the configured threshold, primarily due to the client's weak signal strength.
- **Network Issues** – This classifier tracks the user minutes that client's predicted throughput is below the configured threshold, primarily by the capacity of wired network.

Back to SLE list on page 21

Capacity

This SLE metric tracks the user minutes that a client experiences bad capacity. This metric tracks the per-user available channel capacity and fires off classifiers when the available capacity drops below the specified SLE threshold.

Capacity Classifiers

- **WiFi interference** - This classifier tracks the number of user minutes that a client experiences low capacity that can be attributed to interference.
- **Non-WiFi interference** - This classifier tracks the number of user minutes that a client experiences low capacity that can be attributed to interference.
- **Client Count** - This classifier tracks the number of user minutes that a client experiences low capacity that can be attributed to the number of attached clients.
- **Client Usage** - This classifier tracks the number of user minutes that a client experiences low capacity that can be attributed to client load.

Back to SLE list on page 21

AP Health

This SLE metric is the percentage of time the APs have been operational without losing connectivity to the cloud or rebooting.

AP Health Classifiers

- **Low Power** – This classifier is triggered when the AP is receiving insufficient power from its POE connection.
- **AP Disconnected** – This classifier is triggered by any one of the following sub-classifiers:
  - Switch Down
  - Site Down – This sub-classifier is triggered when all APs on your site are unreachable.
  - AP Unreachable – This sub-classifier is triggered when your AP loses cloud connectivity.
  - AP Reboot
-
- **Ethernet** – The Ethernet classifier contains two sub-classifiers:
  - Speed Mismatch
  - Ethernet Errors

Back to SLE list on page 21

Wired SLEs

IN THIS SECTION

Overview

The wired SLE dashboards display the percentage of time that the SLE metrics met the specified service level expectation goal within a specific time range. These metrics are categorized into classifiers and sub-classifiers, which provide additional details to identify the specific causes of failure. With this information, you can easily identify and address the issues affecting the end-user experience.

Our Wired SLEs are divided into 3 Metrics: **Throughput**, **Switch Health**, and **Successful Connects**. Click into these metrics to find detailed Classifiers and Sub-Classifiers which show the exact circumstances of failures experienced by your clients.

Figure 5: Wired SLE Dashboard

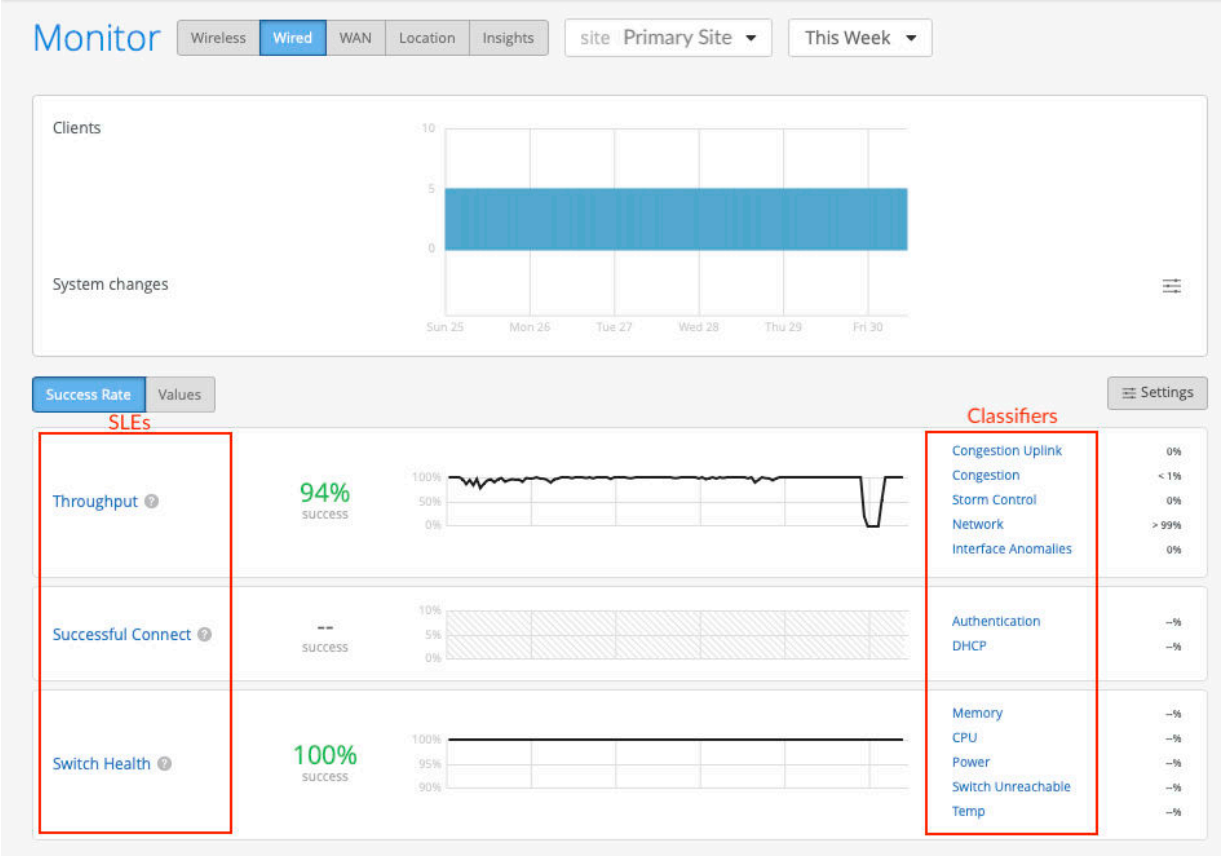

If you click on any of the blue text on the Figure 5 on page 29, you are taken to the Root Cause Analysis Page. In the Figure 6 on page 30 example below, we clicked on the **Throughput** metric. Notice that the root cause analysis page displays a lot of the same information as the Wired SLE dashboard, but provides more information to help you understand how the SLE metrics are influenced by the classifiers.

Figure 6: Wired Throughput Root Cause Analysis

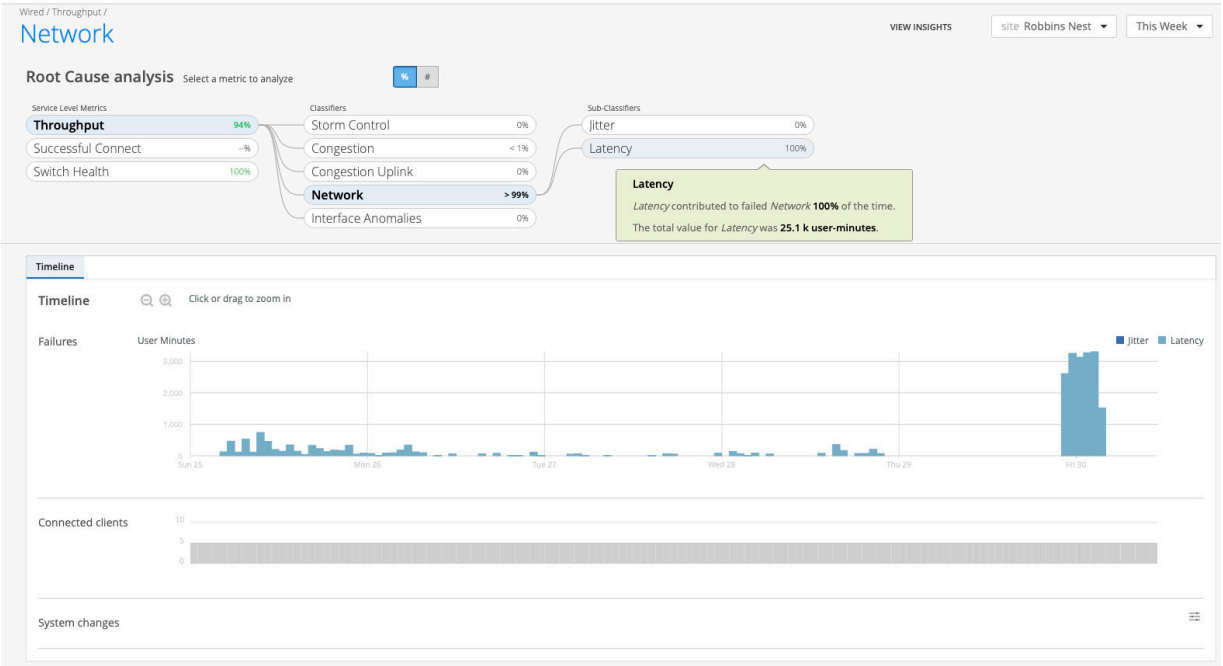

Wired SLE Metrics and Classifiers

This section provides more detail about the metrics and classifiers available for the Wired SLE.

- "Throughput" on page 30
- "Switch Health" on page 32
- "Successful Connects" on page 32

Throughput SLE

The Throughput SLE can help you recognize the need for more wired bandwidth on your site. Many factors can affect the throughput value of your network, including MTU mismatches, bad cables, and devices negotiating at the wrong speed. Using the Throughput SLE, you can proactively assess when your network needs a higher bandwidth to function properly.

Classifiers

The Throughput SLE contains 5 Separate Classifiers:

Congestion - Congestion 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 (TxDrops). We use a formula that takes into account the following 3 ratios to determine if there is a 'bad user minute' due to congestion:

- TxDrops to TxPackets (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)

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 an 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.

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 admin, you can set a maximum transmission unit (MTU) value for each interface. The default value for Gigabit Ethernet interfaces is 1514. To support jumbo frames, you need to configure an MTU value of 9216, which is the upper limit for jumbo frames on a routed 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 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 auto-negotiation failure, duplex conflicts, or user misconfiguration of device settings. Moreover, older devices may not be able 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.

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 levels. These thresholds are known as storm control levels or storm control bandwidth. By default, the storm control level is set to 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—This classifier allows you to monitor user minutes when the throughput is lower than expected due to limitations in uplink capacity. 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 locate 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.

"Back to metrics list" on page 30

Switch Health

Switch health is influenced by several factors, including operating temperature, power consumption, CPU, and memory usage. Monitoring switch health is crucial because issues like 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. The Switch Health metric identifies bad user minutes resulting from the following conditions (listed as classifiers):

- **Switch Unreachable**—The switch can't be accessed.
- **Memory**—The memory utilization is above 80 percent.
- **CPU**—The switch CPU usage is above 90 percent.
- **Temp**—The switch operating temperature exceeds the prescribed threshold range, either going above the maximum limit or below the minimum requirement. For information about the operating temperature supported by Juniper switches, refer to the switch hardware guides in Juniper documentation portal.
- **Power**—The switch power consumption is above 90 percent of the available power.

"Back to metrics list" on page 30

Successful Connect

The Successful Connect metric shows if a client successfully connects to the network. It helps assess the impact of connect failures and identify the issues preventing client devices from connecting to the network.

The Successful Connect metric has two classifiers:

- **Authentication**—Each time a client authenticates, a client event is generated. These could either be successful events or failure events. This classifier helps you identify issues that caused authentication failures. Here's a list of possible reasons for a dot1x 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.
    - There's a routing issue between the switch and the RADIUS server.
    - The RADIUS server is down.
  - If all switch ports on all 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 Windows devices, fails to authenticate, it may suggest an issue related to certifications.
- **DHCP**—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.

"Back to metrics list" on page 30

WAN SLEs

IN THIS SECTION

Overview

WAN Service Level Experience (SLE) is centered around the devices and network functions needed to serve the WAN functions in your network. We use the metrics **WAN Edge Health**, **WAN Link Health**, and **Application Health** to derive the percentage ratings shown in the Figure 7 on page 34. We rely on telemetry and health data from the Juniper WAN edge device itself for the metrics listed above. An example of the WAN SLE dashboard is shown below.

Figure 7: WAN SLE Dashboard

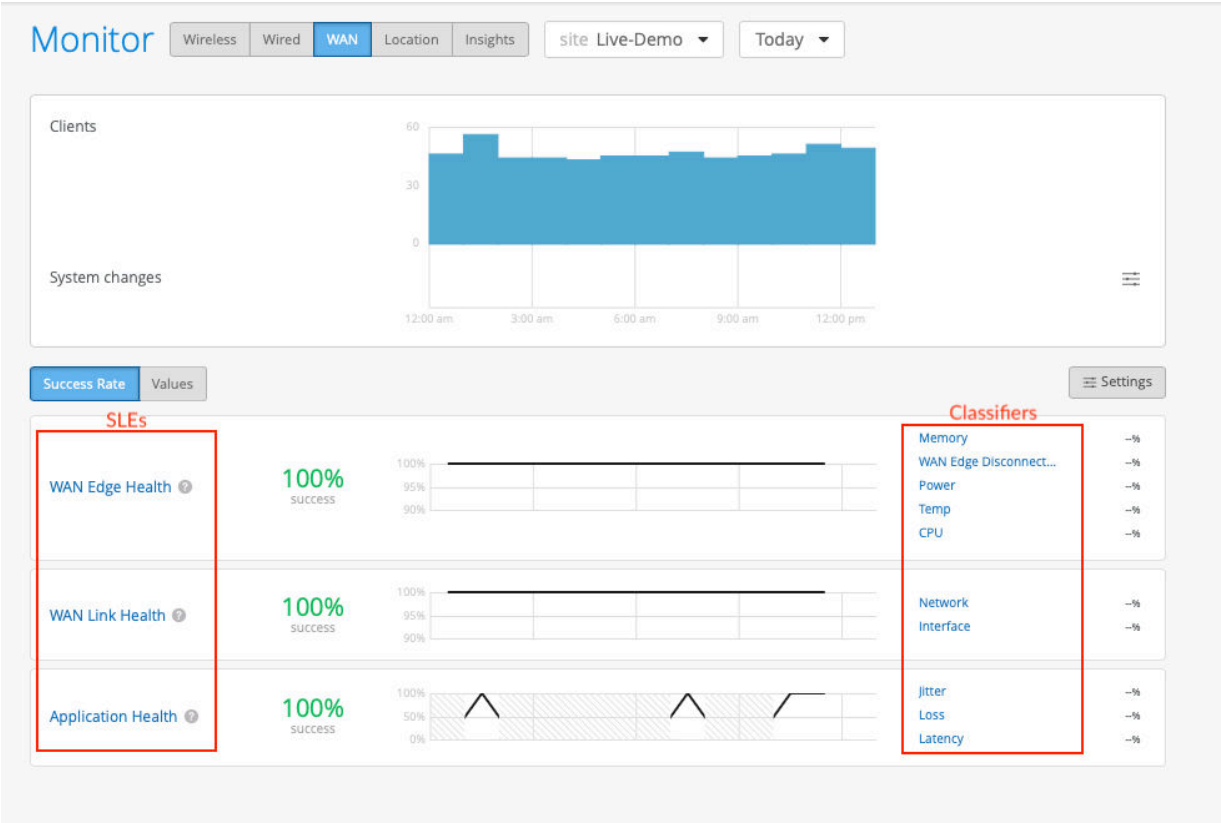

As you can see, we display the metrics on the left and the classifiers for each metric on the right. When you click on any of the blue text in the dashboard, you are taken to the Root Cause analysis page for the link text you clicked. An example of the Root Cause analysis page for the Jitter sub-classifier is shown in Figure 8 on page 35.

Figure 8: WAN SLE Root Cause Analysis Page

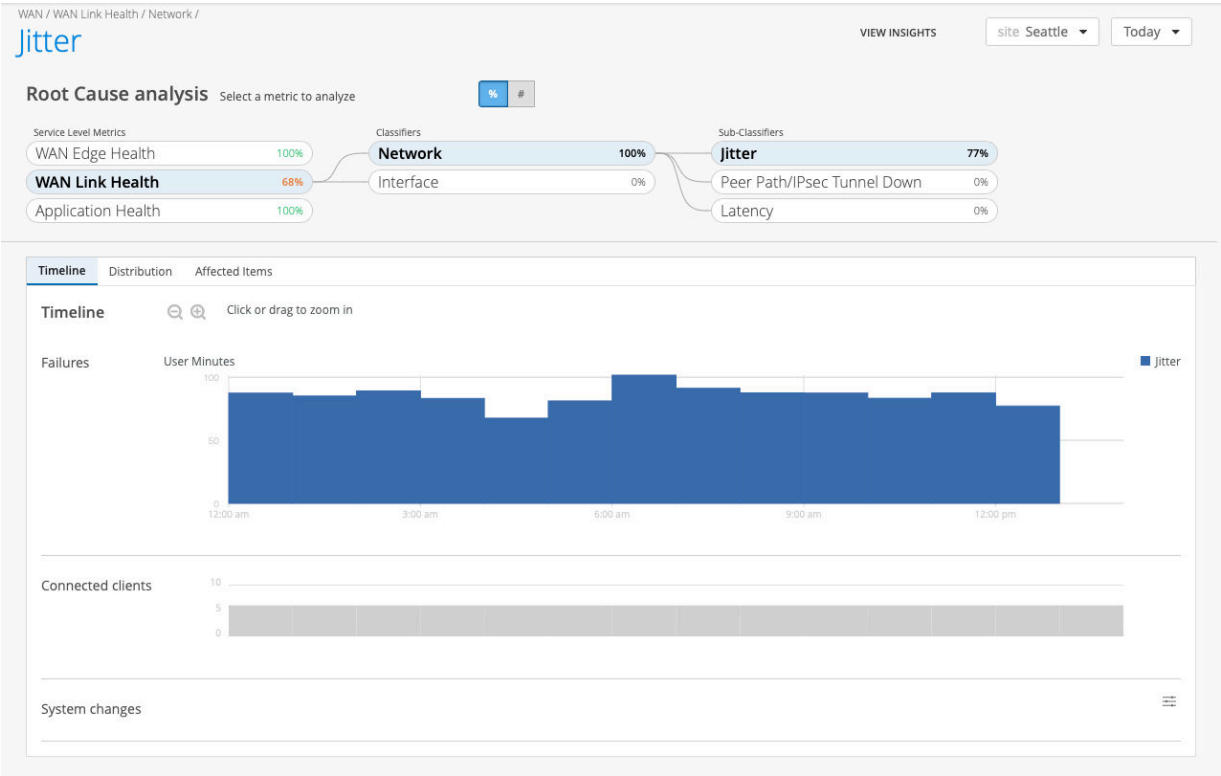

As you can see, a lot of the same information about the WAN link health metric is shown in the root cause analysis. The difference is that we provide more detail including sub-classifiers, if they are available.

The following sections provide detailed information about each of the WAN SLE metrics, their associated classifiers, and sub-classifiers (if they are available.)

WAN Edge Health

WAN edge health is the percentage of time when the health or performance of the WAN edge device was not optimal. This affects the device's ability to pass traffic, thus directly affecting any clients connected to the device.

WAN Edge Health Classifiers

- - **Memory** – This classifier is set when the WAN Edge memory utilization is above 80 percent.
  - **WAN Edge Disconnected** – This classifier is set when the WAN Edge is disconnected from the Mist Cloud.

- **Power** – This classifier is set when power consumption is above 90 percent of the available power.
- **Temperature** – This classifier is set when the WAN edge operating temperature exceeds the prescribed threshold range, either going above the maximum limit or below the minimum requirement.
  - **CPU –** This sub-classifier is set when the CPU temperature exceeds the prescribed threshold range.
  - **Chassis –** This sub-classifier is set when the chassis temperature exceeds the prescribed threshold range.
- **CPU** – This classifier is set when the WAN Edge when CPU utilization is above 90 percent. When the CPU utilization spikes on a Juniper WAN edge device, downstream devices can lose their connectivity, thus directly impacting clients and preventing them from being able to pass traffic.
  - **Data Plane** – This sub-classifier is set when the Data Plane CPU utilization is above 90 percent.
  - **Control Plane** – This sub-classifier is set when the Control Plane CPU utilization is above 90 percent.

WAN Link Health

WAN link health is the percentage of time when the WAN link health was bad. Bad link health affects the device's ability to pass traffic, thus directly affecting any clients using that link.

WAN Link Health Classifiers

- **Network** – This classifier tracks network issues that contribute to WAN link health. The Network classifier has 3 sub-classifiers:
  - **IPSec Tunnel Down** – This sub-classifier is triggered if the SRX WAN Edge detects that one of the Overlay IPSec tunnels is down.
  - **Latency** – This sub-classifier is activated when WAN link traffic shows latency. Latency is calculated using the average value of round trip time (RTT) for traffic over a period of time.
  - **Jitter** – This sub-classifier is activated if the WAN link experiences by jitter. Jitter is calculated using the variation (standard deviation) of RTT within a 5 to 10 minute time period for a particular WAN link. We compare the calculated value with the average deviation of RTT over a day or a week.

- **Interface** – This classifier tracks interface issues with the WAN link health. The Interface classifier has 3 sub-classifiers:
  - **Cable Issues** – This sub-classifier is triggered when we detect faulty cables in the network.
  - **Congestion** – This sub-classifier is activated when congestion is negatively affecting WAN link health. Congestion measures the number of output packet drops. When packets come into an interface, they are queued in a buffer. When the buffer becomes full it will drop packets (TxDrops).
  - **VPN** – This classifier is activated when there is a performance issue with the VPN.

Application Health

Application health is the percentage of time when the monitored applications are performing poorly thus resulting in bad user experience.

Application Health Classifiers

The Application health classifier has 3 sub-classifiers:

- **Jitter** – This classifier is triggered when application health is affected by jitter. Jitter is calculated using the variation (standard deviation) of RTT within a 5 to 10 minute time period for a particular WAN edge or link. We compare the calculated value with the average deviation of RTT over a day or a week.
- **Loss** – This classifier is set when Application Health is affected by packet loss. The loss value is calculated based on the number of lost packets over a period of time.
- **Latency** – This classifier is set when Application Health is affected by latency. The latency value is calculated based on the average value of RTT over a period of time.

PART 3

Alerts

Alert Configuration

SUMMARY

Learn about alert monitoring in the Mist dashboard.

IN THIS SECTION

Alerts Overview

IN THIS SECTION

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 infrastucture alerts. For example, if a Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), or RADIUS server is unreachable, many clients could be affected. 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.

Marvis alerts are raised by the Mist Predictive Analytics and Correlation Engine (PACE) and signify those events that Marvis tracks. For example, if an access point (AP) regularly fails health checks, Marvis will notice it and track it.

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. If a client connects to a rogue AP, that could be even worse.

Within each type, we rank each alert by severity:

Table 2: Alert Severity

Severity	Indicator	Recommended Action
Critical	Red dot	Requires immediate attention
Warning	Orange dot	Continue monitoring in case it continues
Informational	Blue dot	No action is required.

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

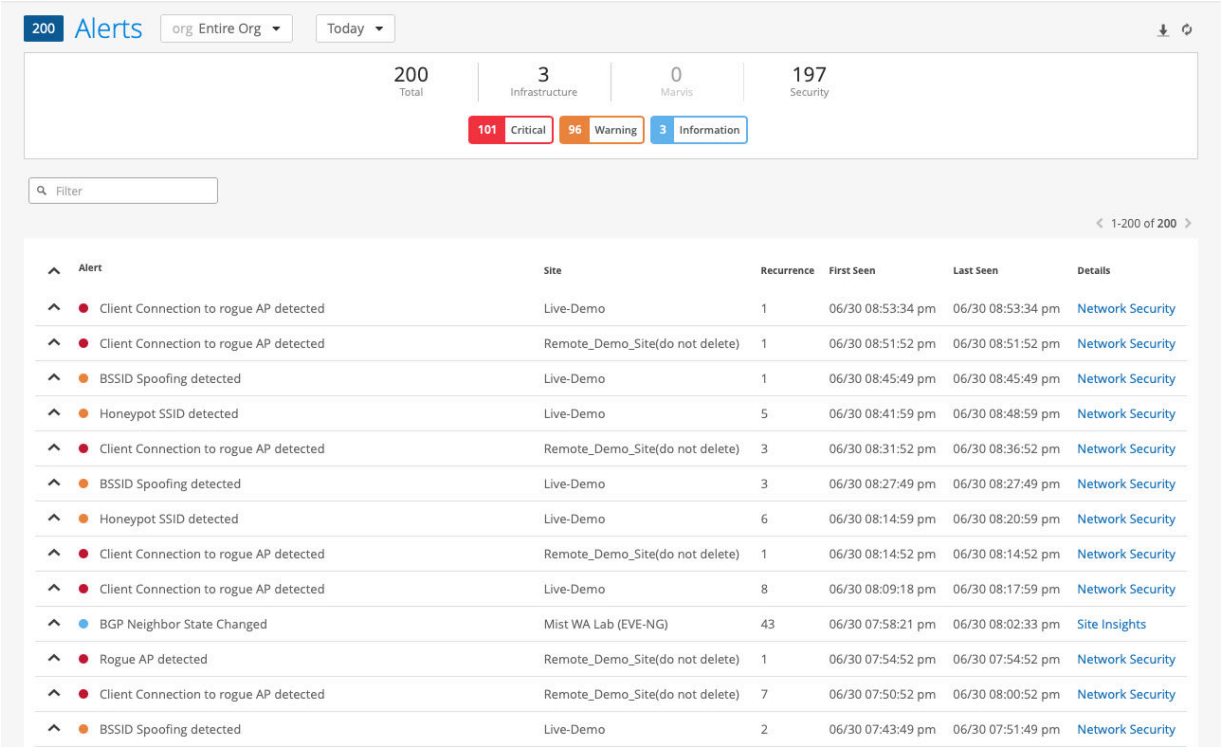

With the drop-down lists at the top of the Alerts dashboard, you can filter the dashboard display by:

- Site—You can choose to see alerts from your entire organization or any site for which you have view access. 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 to 7 days.

Alert Configuration

When you click **Alerts Configuration**, 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 the check box 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 9 on page 42 page below.

Figure 9: Alert Configuration

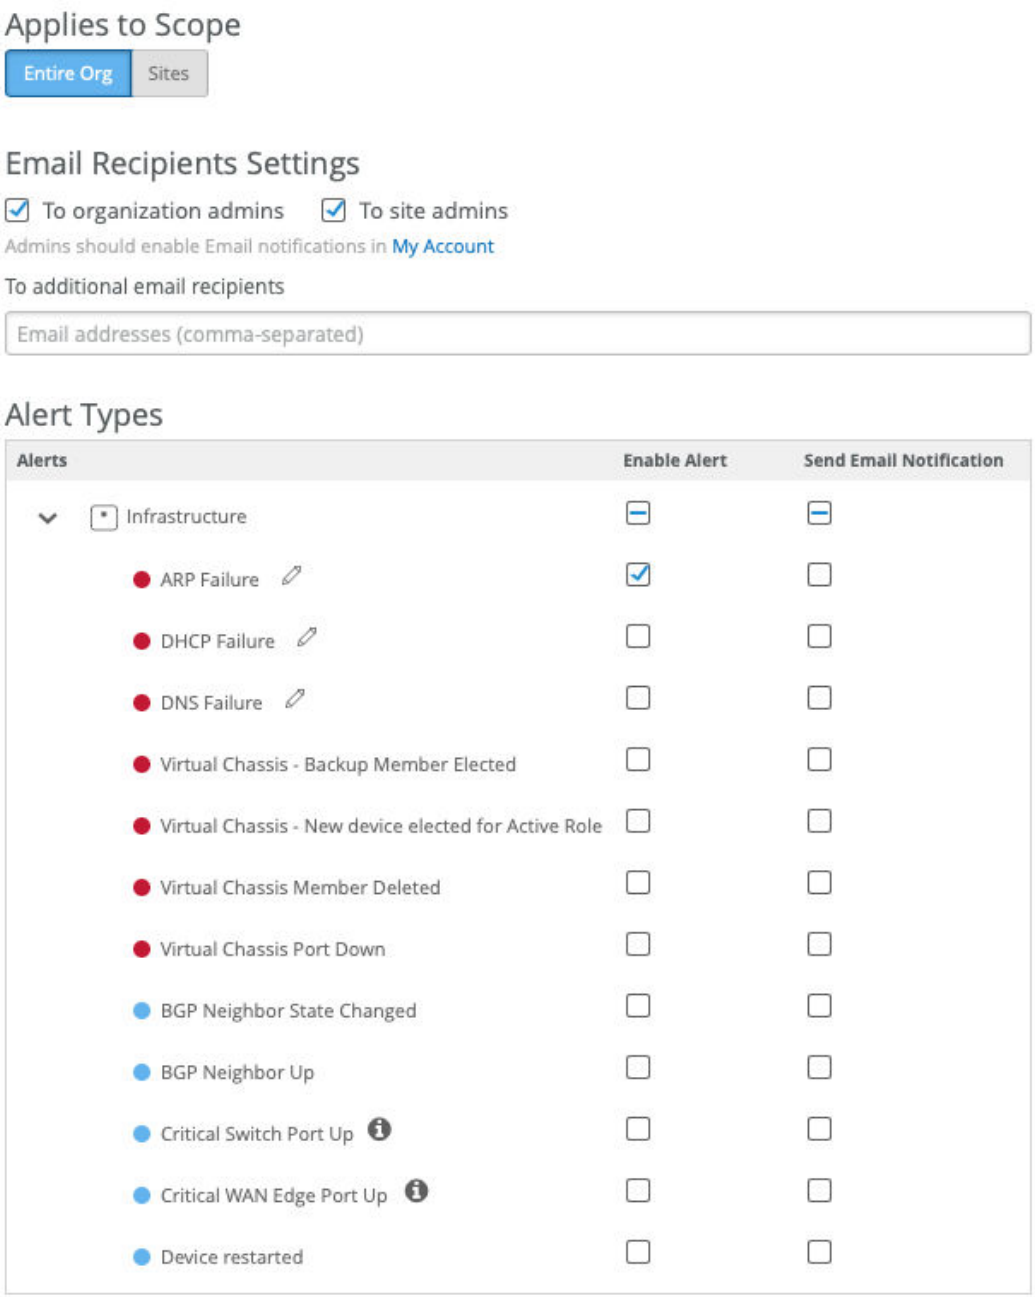

on page 42

In the Figure 9 on page 42 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 indicating 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 indicating that you need to perform additional configuration to enable these alerts.

The two special icons described above only appear 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:

Mist Alert Types

IN THIS SECTION

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 ongoing network and device issues. In the **Monitor > Alerts** dashboard, you can see three different types of alerts: Infrastructure on page 44, Marvis on page 46, and Security on page 48.

Mist categorizes alerts that potentially affect a large number of clients as infrastucture alerts. For example, if a DNS, DHCP, or RADIUS server is unreachable, many clients could be affected. 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.

Marvis alerts are raised by the PACE and signify those events that Marvis tracks. For example, if an AP regularly fails health checks, Marvis will notice it and track it.

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

Table 3: 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	
Informational	DHCP Recovered	X
Informational	DNS Recovered	X

Table 3: Infrastructure Alerts by Severity *(Continued)*

Severity	Alert Name	API Only
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	
Warning	Switch PEM Alarm	
Warning	Switch PoE Alarm	

Table 3: Infrastructure Alerts by Severity *(Continued)*

Severity	Alert Name	API Only
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. For example, if an AP regularly fails health checks, Marvis will notice it and track it.

Table 4: 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 activity 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. Mist lists all security alerts except those that relate to IDP or URL filtering on the Monitor > Alerts page. You can find IDP and URL filtering events and their severity listed on the **Site > WAN Edge > Secure WAN Edge IDP/URL Events** page.

Table 5: 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

Network Monitoring with Juniper Mist | 2

Alert Configuration | 39