

IMPLEMENTING BRANCH NETWORKS FOR AI-DRIVEN ENTERPRISE CUSTOMERS

Using Mist AI to Create, Deploy and Manage a Full Network Stack

Challenge

- Legacy networks focus on connectivity, but lack insight into experience (“up” is not the same as “good”)
- Deployment and configuration are cumbersome—too many manual operations
- With distributed enterprises you must manage branches, HQ(s), DC(s), cloud(s); as a Managed Service Provider (MSP) you must do this for multiple organizations
- These domains are often siloed networks with no shared information
- Existing solutions lack end-to-end visibility and control to manage complexity

Solution

- Mist AI controlling everything internally (Wi-Fi, wired) to sites as well as intersite (WAN)
- AI-driven solution that is born in the cloud
- Templated for easily duplicable and multisite installations and updates
- Claim codes and simplified shared key usage without NAC overhead
- MSP Dashboard allows you to manage an entire customer estate
- SLEs for insight into real-time user experience anywhere on the network

Benefits

- Simple, fast, accurate
- Secure (zero trust) and scalable
- AIOps to reliably maintain/update the entire enterprise: all elements at all sites
- Bandwidth optimization and reduction
- Increased operational efficiencies

Introduction

This solution guide provides unique insight into Juniper’s AI-driven Enterprise (AIDE) via the implementation of wireless, wired and wide area solutions in a branch office. Juniper Mist Cloud is used to operate [Session Smart Routers \(SSR\)](#), [EX switches](#), and [Mist access points](#).

By illustrating a walkthrough of deploying a branch for your customer, this guide highlights the key components of how a full stack AIDE (router, switch, and access point) is built, and showcases the unique advantages of using the AIDE to build an experience-first, [client-to-cloud](#) network for a distributed enterprise.

The process of planning, deploying and managing the AIDE has become easier than ever before, and this guide provides the highlights, along with public resource links such as documentation and videos for more information. As a Managed Service Provider (MSP), you can use the [Juniper MSP Dashboard](#) to manage your entire customer estate.

Furthermore, you can see firsthand how to perform the tasks in this document by setting up an account at manage.mist.com. There are tutorials available and your account representative can help you get started.

Please contact your Juniper representative to set up a guided demonstration of all the features covered in this document. To learn more about how to become a provider in Juniper’s Managed Services program, consult the [Juniper Unified Managed Services](#) page.

The Branch Network in the Distributed Enterprise

In the example shown here (*Figure 1*), there will be three network entities created: a router, a switch, and an access point:

- The router is an SSR120 Session Smart Router (SSR) (interchangeable with all SSR devices)
- The switch is an EX 2300 (interchangeable with a 3400, 4100, or 4300 series)
- The access point is an AP41 (interchangeable with all APs)

Deploying a full stack branch—including the router, switch and access point—with [Mist AI](#) is very straightforward and intuitive.

Topology

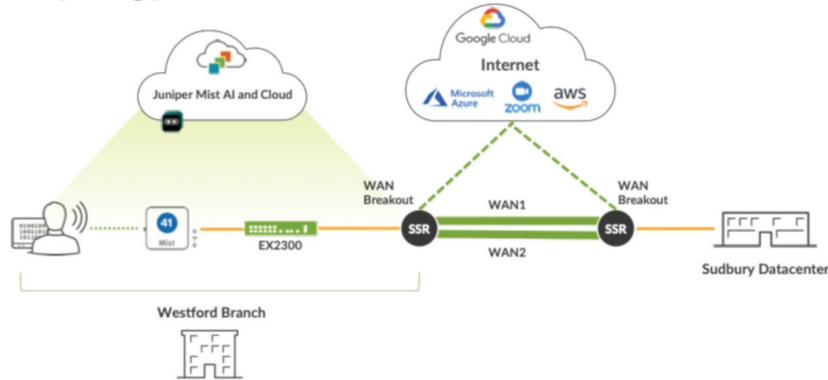


Figure 1: Full Stack Deployment Topology

Figure 1 shows a topology diagram of the full stack deployment.

In this example, there will be two WAN links connected to the datacenter and one for Internet Breakout, which will consume cloud services. The data center is already configured, and only the branch is being deployed. The Mist access points, EX switches, and SSRs are all managed by Mist AI.

There is full documentation for the tasks discussed in this guide:

- For the SSR, see [Guided Setup for Juniper Mist WAN Assurance](#)
- For the EX and Mist AP, see [Configure and Manage the EX Switch and Mist AP in the Juniper Mist Cloud](#)

These will all be managed by Mist AI, which allows you to deploy, operate and maintain all customers from a single console.

Within this console, Wi-Fi Assurance, Wired Assurance and WAN Assurance simplify every part of administration and deployment as well as with root cause isolation and correction, anomaly detection, insights from client to cloud, and AI integration for the help desk.

This discussion is illustrated in deployment stages (Figure 2).

Industry-wide, the stages are defined as follows:

- Day 0: Planning and Design
- Day 1: Zero Touch Deployment
- Day 2: AI-driven Operation and Benefits
- Day 2+: Ongoing Maximization of the User Experience

Note: See the [Resources](#) section below for other helpful documentation and content assets.

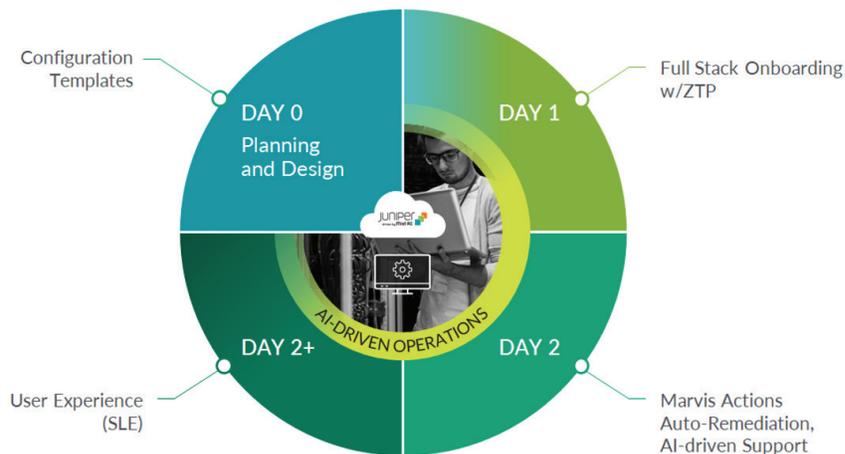


Figure 2: Deployment Stages

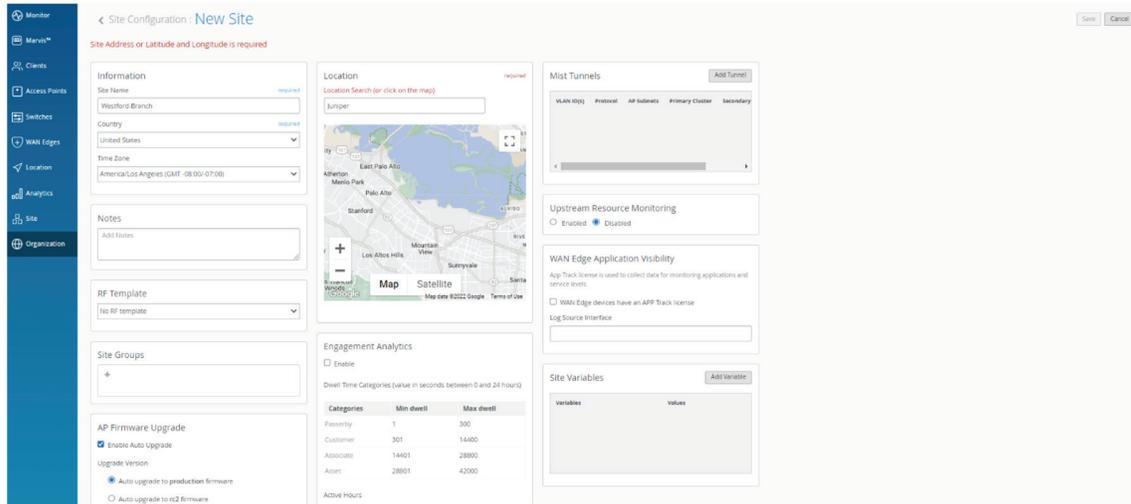


Figure 3: Selecting a Location for the Site

Day 0: Planning and Design

When you log into manage.mist.com, you will be presented (assuming you have been granted access) with the **Managed Service Provider (MSP) Dashboard** and the organizations within your estate. The process of setting up organizations, applications and networks begins by defining the sites within the organization.

Setting Up Organizations, Applications and Networks

To start, click the “Create Organization” button in the upper right. Once initialized, the new organization will appear in the list below. Click into the organization to continue with configuration.

In the menu, under *Organizations*, you can set up your administrators, and invite new ones if need be. You can give the administrators different roles, up to super user authority.

Site Configuration

To create a site under *Organization*, select *Site Configuration*. After giving the site a name, you select a location for the site (Figure 3).

Using Site Variables

For automation and scaling purposes, you can assign variables for a variety of items. These variables can be called when claiming devices later in the process.

For instance, in Figure 4, you can see this site variable being entered in a site definition. Thus, you can configure many sites with only a few templates.

For more detail on how the use of site variables facilitates site creation under Mist, see the [Guided Setup for Mist WAN Assurance](#).

Setting Up Networks and Applications

In configuring the SSR, you determine the network layout. This includes specifying access groups—the users and devices that will access the network and will be made reachable. Access groups define the entities that will access certain resources (applications or services).

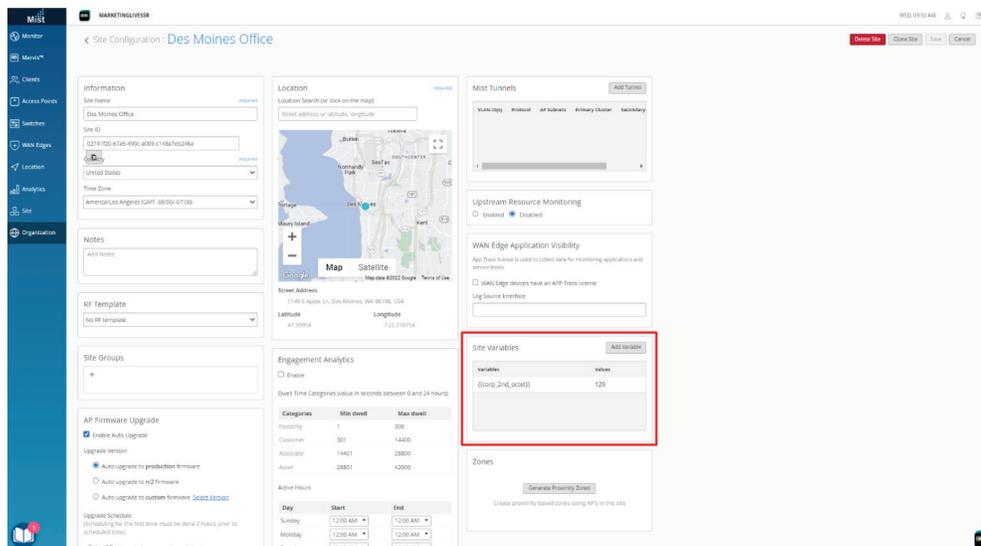


Figure 4: Site Variable Entry

NAME	SUBNET	VLAN ID	USERS	ADVERTISED VIA OVERLAY	STATIC SOURCE NAT	DESTINATION NAT	SOURCE NAT POOL
branch1	10.0.{{branchnumber}}.0/24	--	--				
branch-security	10.255.{{corp_2nd_octet}}.0/24	255	--	✓			
corp-branch	11.0.1.0/24	--	--	✓			
corp-dc	11.0.0.0/24	--	--	✓			
corp-user	192.168.{{corp_2nd_octet}}.0/24	--	--	✓			
dc-lan	10.0.0.0/16	--	--				
guest	192.168.1.0/24	254	--				
home	192.168.130.0/24	--	--				
lan	10.0.0.0/16	--	--				
westford-lan	10.0.1.0/24	--	--				

Figure 5: Networks for An Organization

This is achieved with a zero-trust security model: none of the entities (devices or users) in the access group can access any applications by default. There is no access unless it is explicitly granted.

Networks

You thus set up *Networks* (Figure 5), which in SSR terminology specifies the *tenants* (who) that access *services* (what).

This is where operators can divide the branch LAN into multiple VLANs, each with its own access group. For example, a drug store chain might have one VLAN used for point of sale, one VLAN used for ATM, another VLAN for the pharmacy, etc.

Applications

Once the *Networks* are globally defined, the next step is to define *Applications*. While networks include tenants, applications are the services that those tenants utilize. Without a defined association between networks and applications, traffic is not permitted on the SD-WAN—this is the deny-by-default nature of SD-WAN driven by Mist AI.

Thus, you must define all the applications to which you want to route. Applications can be selected off a list or via an application category, or defined by a customized approach.

Some applications are predefined for operator selection: Office 365 is a common example. Operators might want to have special routing when traffic is going to a Microsoft cloud versus (for instance) to a non-business application such as a social media site for personal use. Depending on corporate policies, operators might want to have special routing or block that type of traffic. Similarly, it would be common to prioritize conferencing and business over all personal (but permitted) activities.

Applications and services will be set up in the routers as access control lists (ACLs), allowing networks to access those applications under certain conditions. Adjustments can be made as operators notice patterns in the routing that could be optimized. For instance, if video conferencing is slow it may need higher prioritization.

Under *Application Categories*, for instance, you might see items such as conferencing, file sharing, or financial (Figure 6).

NAME	TYPE	HEURISTIC TYPE
SVC1	Apps	Custom
corp-internet	Custom	Default
corp-network	Custom	Default
corp-dc-lan-network	Custom	Default
corp-branch-lan-network	Custom	Default
M365	Apps	Default
branch-lan	Custom	Default
guest-internet	Custom	Custom
isp-int	Custom	Custom
internet	Custom	Default
zoom	Apps	Default
Youtube	Apps	Default

Figure 6: Sample Application Categories

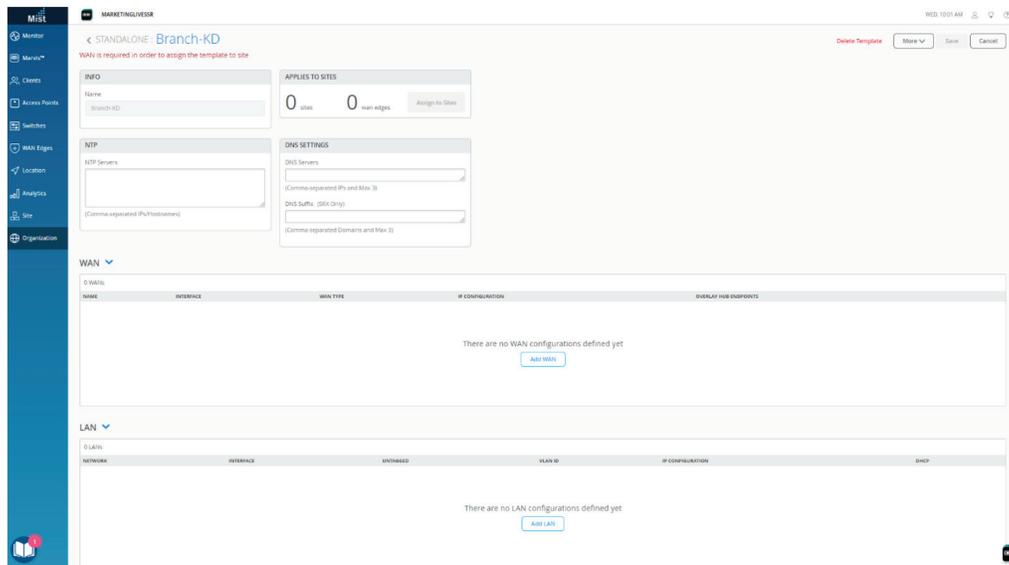


Figure 7: WAN Edge Templates

Using these categories, operators can select an entire section or a set of applications; special treatment defined here might include blocking of adult web sites.

Note: SSRs recognize over 7,000 applications using over 230,000 identifiers.

WAN Edge Templates

Deploying a branch network continues by creating templates for the network elements. The SSR at the branch edge will be defined using WAN Edge Templates (Figure 7).

Note: You can create a new template or use an existing one. If you're going to connect it to a hub (data center), you'll make it a spoke (branch).

A key advantage of templates is that if a supported organization has many sites (even into the thousands) and a change needs to be made to all (or a large subset) of them, that change can be made in one spot and applied to the other sites.

In these WAN Edge Templates you input:

- *General Information* such as the name of the template and which sites it applies to
- *Servers* for NTP and DNS
- *WAN Interfaces*: names and descriptions; this includes whether DHCP, NAT, or traffic shaping is enabled
- *LAN Interfaces*: names and descriptions; this is where you assign the Networks that were previously created, and it's also where you create VLANs for access groups or for device types such as IoT
- *Traffic Steering* which defines policies for steering traffic through selected interfaces
- *Application Policies* which define access for the traffic steering as well as branch security policies for IDP
- *Routing* for links between branches and hubs, where traditional routing protocols or static routes might be used

All of the above is discussed in more detail in the [Guided Setup for Juniper Mist WAN Assurance](#).

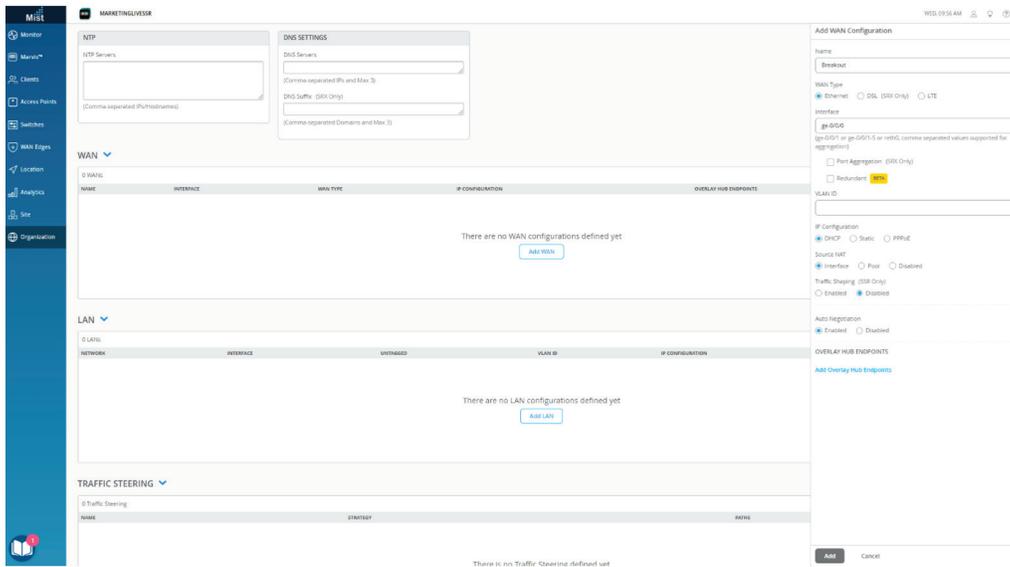


Figure 8: Sample WAN Interface Definition

Defining Interfaces

A sample WAN Interface definition is shown in Figure 8.

For example, consider a common case where you would have one WAN Interface for Internet Breakout and two others for redundant and load-balanced connections to a data center. For each of the WAN links, you can define the physical interface, the type of WAN (i.e., broadband or LTE), the IP

configuration, and the overlay hub endpoints (for instance, an SSR in a data center) for the interfaces.

Note the Overlay Hub Endpoints selection in Figure 8; this is where you tell the spoke (branch) about the hub (data center) endpoints.

All three of the defined WAN interfaces discussed here are shown in Figure 9.

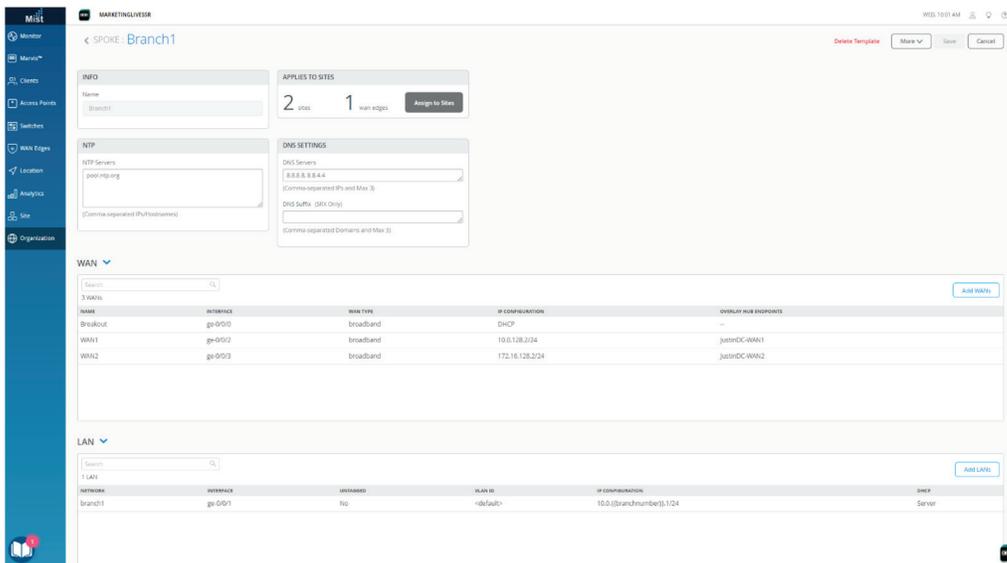


Figure 9: WAN Interfaces for Internet and for Data Center Connections

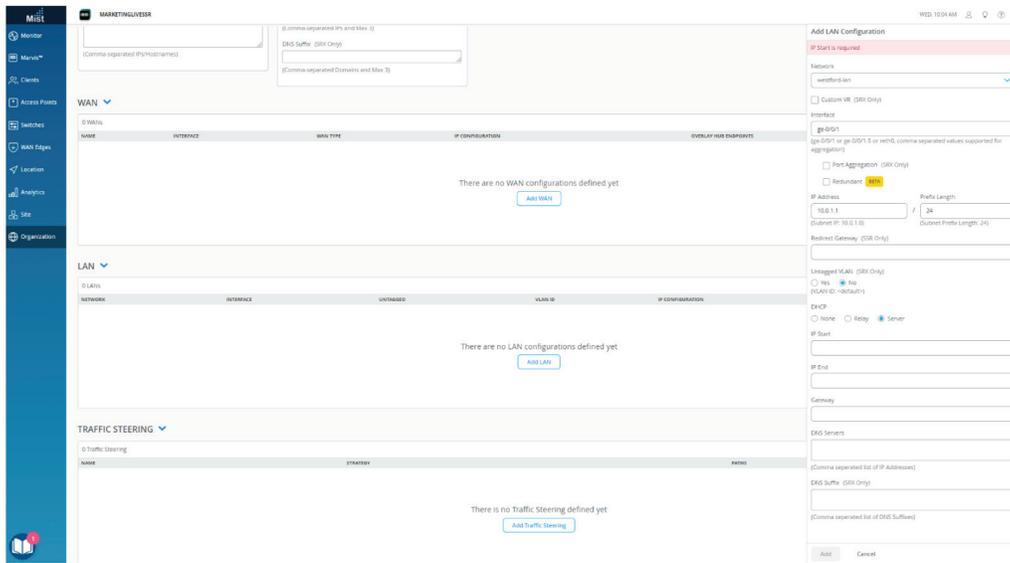


Figure 10: Sample LAN Interface Definition

Of course, the SSR also connects to the LAN; sample LAN Interface definition is shown in Figure 10.

Traffic Steering and Application Policies

Once all the WAN and LAN interfaces are defined, traffic steering can be set up. You can set up multiple traffic steering policies. The one shown in Figure 11 is for Internet Breakout traffic.

For any traffic steering policy, you can include several paths to be included in that policy, as well as the strategies for utilizing those paths.

There are different “strategies” that you may choose; these are:

- *Ordered*: Start with a specified path and failover to backup path(s) when needed
- *Weighted*: Distribute traffic across links according to a weighted bias, as determined by a cost that you input
- *ECMP*: Equal-cost multipath; load balance traffic equally across multiple paths

In the case of Internet Breakout, *Ordered* is often selected. As this is breakout traffic, there is no overlay connection to a data center.

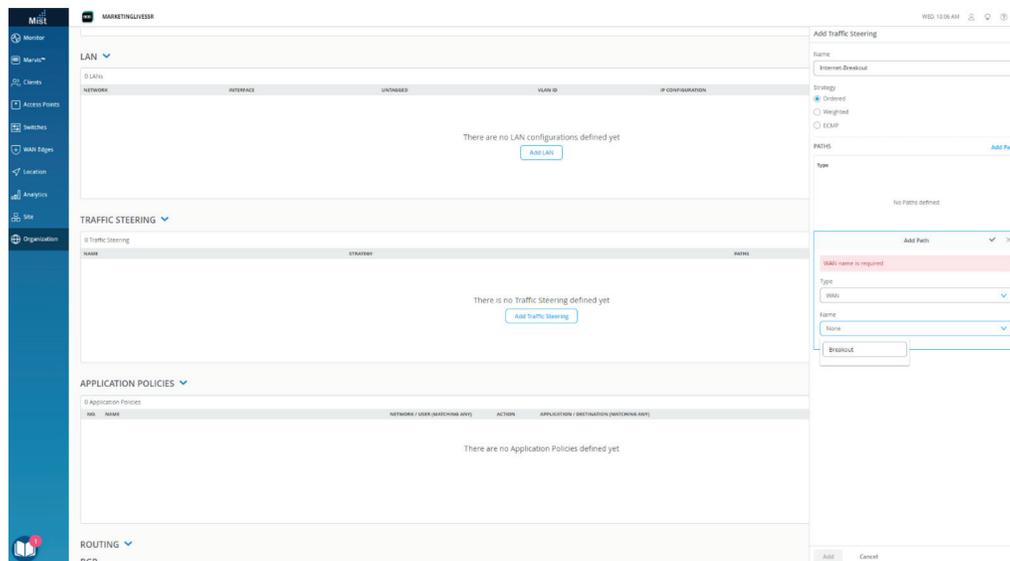


Figure 11: Traffic Steering Policy Setup for Internet Breakout

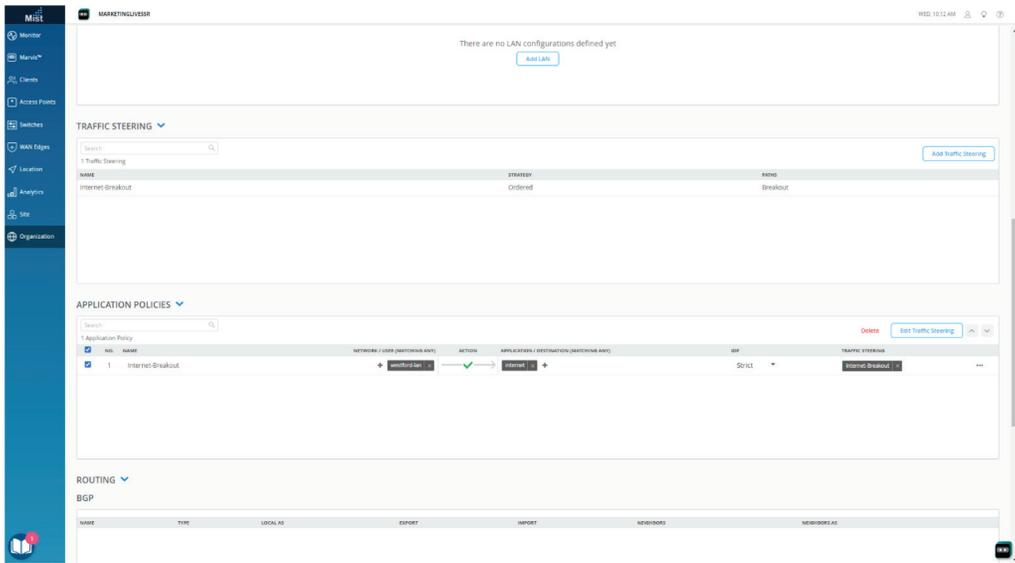


Figure 12: Application Policy Definition and Application to Traffic Steering Policy

Having set up these traffic steering policies, you can set up application policies, which define “who” (which access group) is going to use certain traffic steering policies. For instance, Figure 12 shows an application policy allowing the LAN to access the Internet.

The selected traffic steering policy tells the network what to do when traffic matches this application policy. In this case, the traffic will be referred to the Internet application policy, which forwards traffic out of the Internet Breakout interface.

For traffic destined to applications in the data center, you would set up a different traffic steering policy for data center backhaul (Figure 13).

In this example, the LAN access group can access corporate network applications and IP camera (ipcams) applications, and will use the DC Backhaul application policy.

You could also set up static or BGP routing if needed.

After saving this, the branch WAN template is defined, and you can move on to defining templates for the switches and access points.

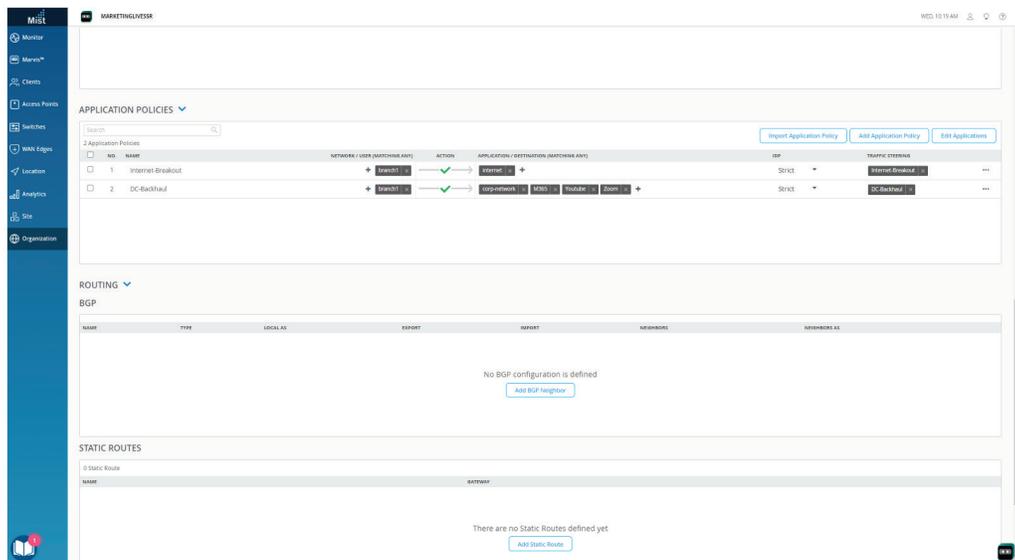


Figure 13: Application Policy Definitions for Data Center Traffic and Application Steering

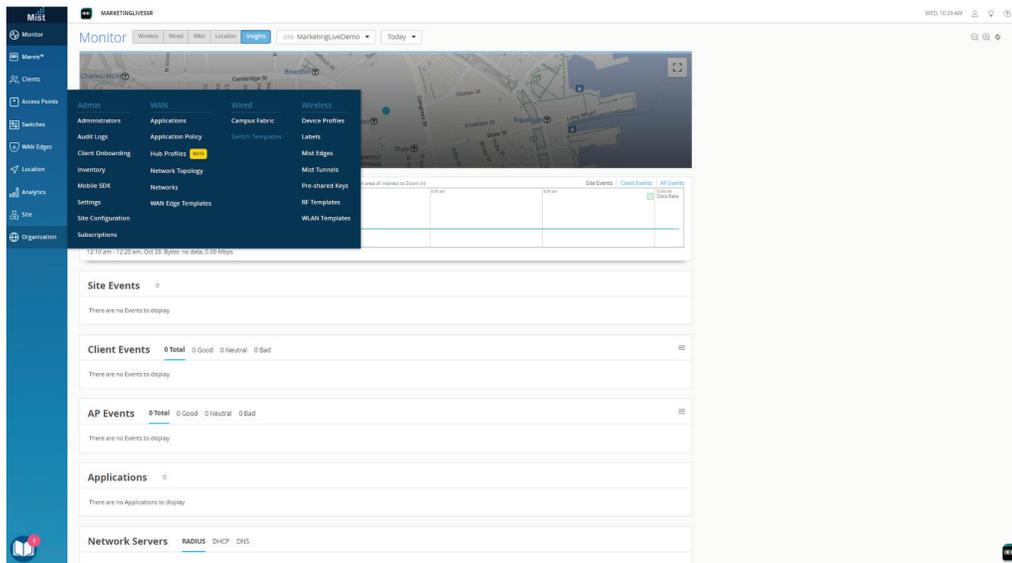


Figure 14: Selecting Switch Template Under Wired

Switch Templates

Setting up new switch templates begins by going to *Organization*, then selecting *Switch Template* under *Wired* (Figure 14).

You then have a blank branch template to begin filling in switching parameters (Figure 15).

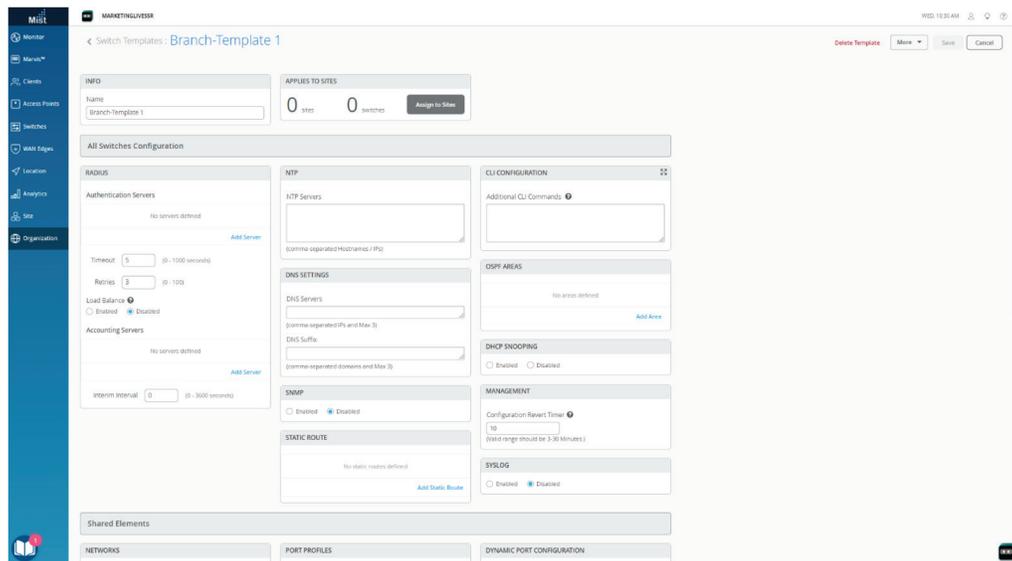


Figure 15: Switch Template

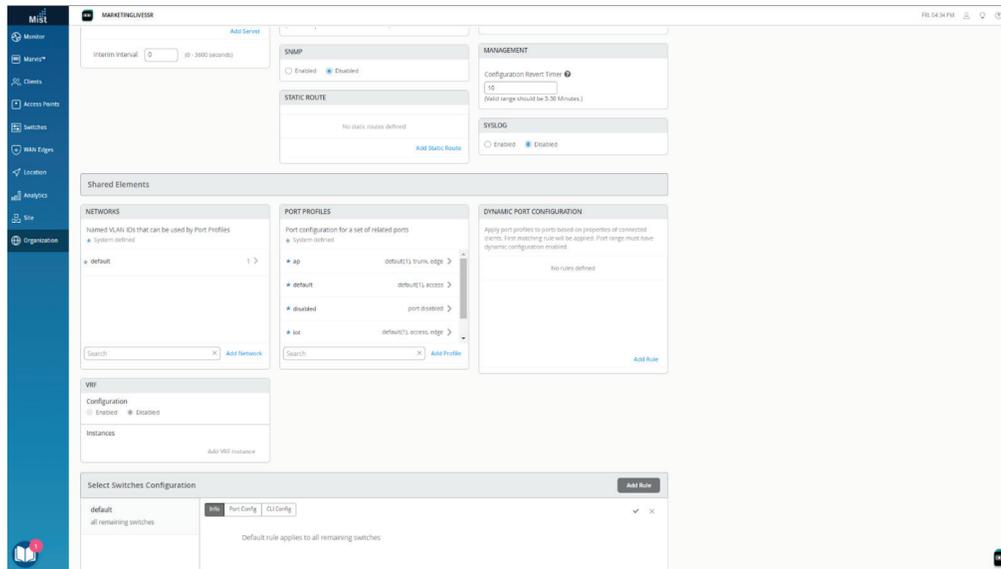


Figure 16: Shared Elements

For instance, you can set up RADIUS servers for authentication or accounting, DNS settings, OSPF areas, static routes, or management. Under Shared Elements (Figure 16), you can create networks, port profiles and dynamic port configuration.

You can set up static or dynamic rules for port configuration. Using dynamic port configuration, you can ensure that the switch recognizes the roles of network elements with certain MAC addresses (such as an access point).

For more information on setting up an EX switch, see the section in the documentation entitled, [Configure and Manage the EX Switch and Mist AP in the Juniper Mist Cloud](#).

WLAN Templates

Select WLAN Templates under Organization/Wireless to see the following template (Figure 17).

Name	Applied to Org	Site	Site Group	Exceptions	WLANs
Branch1-Template	No	Juniper-Westford, Westford-Branch			corp, Guest, Justin
Cloned-Template	No				EDGEADSECTEST, Goddard-tunnel-dot1x, Goddard-Tunnel-Guest, Goddard-tunnel-psk, Jarvis, mist-remote-wifi
Home	No	Justin Time			JustIFI
HomeJM	No				Corp
LD-Inter-Op	No				LD-CBRS-C, LD-CBRS-G, Test-BC
Live Demo WLAN Template	No				Guest_Live_Demo, Live_demo_dfs_not_remove, Live_demo_only, Mist_Corporate, Mist_Just
Remote WFH	No				EDGEADSECTEST, Goddard-tunnel-dot1x, Goddard-Tunnel-Guest, Goddard-tunnel-psk, Jarvis, mist-remote-wifi
Remote_Demo_Donotdelete	No				Remote_Demo_donotdelete
small-branch	No	North Bend Office			Acme-Corp-WLAN
test	No				CORP

Figure 17: WLAN Template List

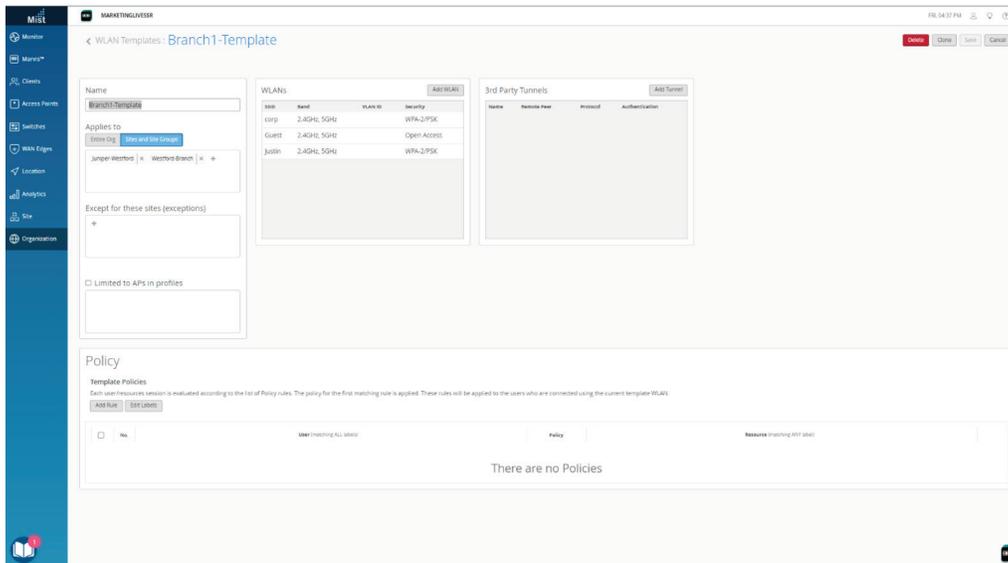


Figure 18: A New WLAN Template

You begin by creating a template and giving it a name (Figure 18).

Under *Applies To*, you can assign the template to a site. But first, you would just add a WLAN (Figure 19) to create a wireless network.

Here, you can assign an SSID, set up Wi-Fi Settings, radio bands, data rates, access policies, rate limiting, and other advanced settings. You can also set up Security (such as WPA) and a guest portal.

You can create multiple SSIDs. For instance, you can create both a corporate network and a guest network under Create WLAN.

For more information on setting up a Mist AP, see the documentation section entitled, [Configure and Manage the EX Switch and Mist AP in the Juniper Mist Cloud](#).

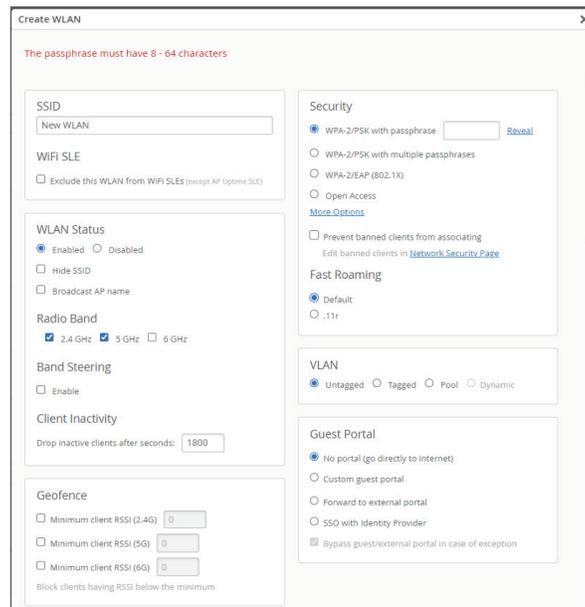


Figure 19: Add WLAN

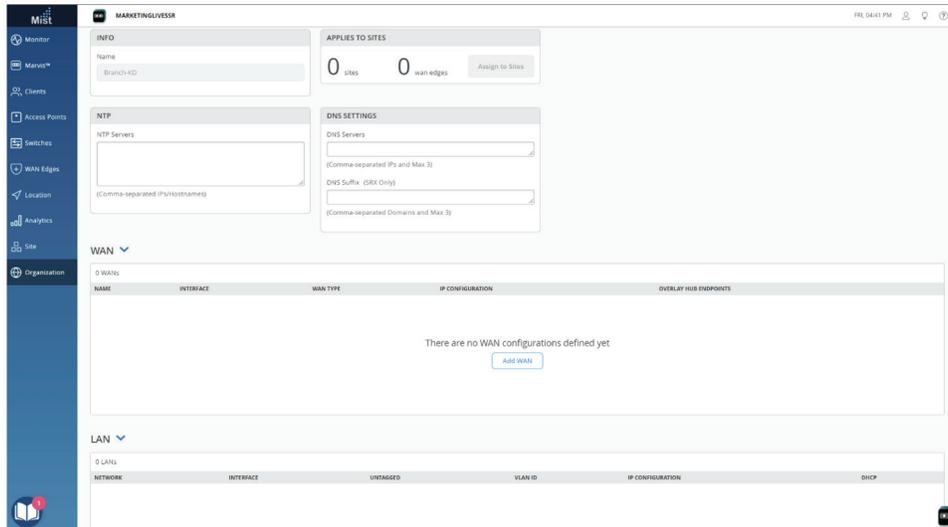


Figure 20: Assigning a Router Template to a Site

Assigning Templates to Sites

Once you've created your sites and completed your templates for each domain, you can then assign the templates to them. For every template, select Assign to Site (Figure 20).

This process is the same for switch templates and access point templates.

Figure 21 shows the access point template being assigned to a branch.

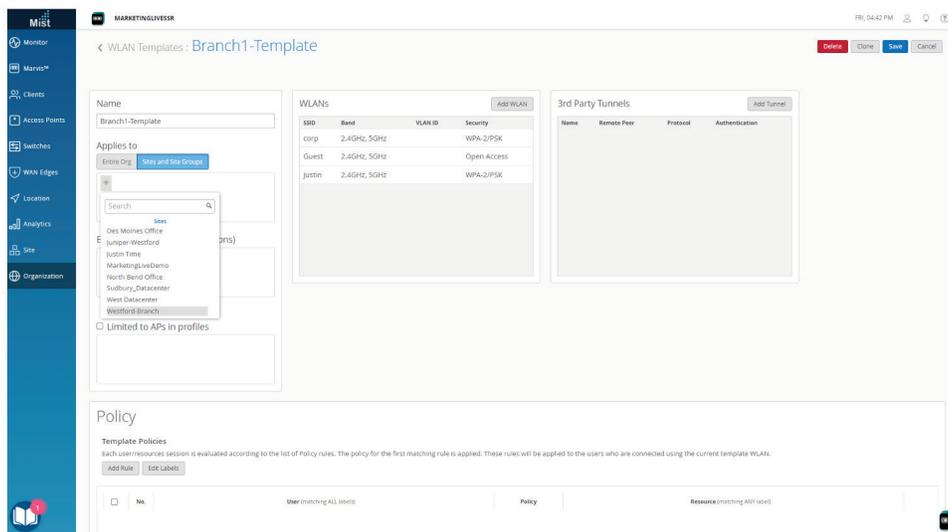


Figure 21: Assigning an Access Point Template to a Site

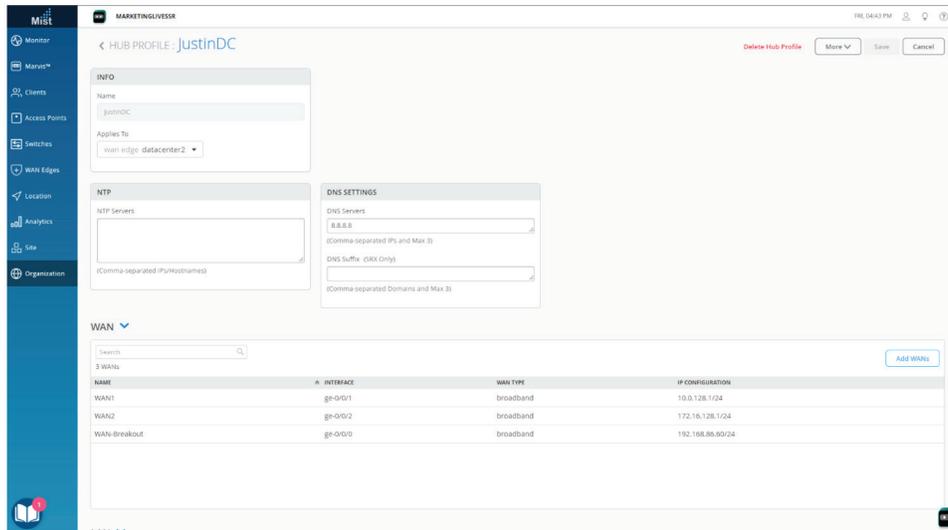


Figure 22: Hub Templates

Hubs and Data Centers

Operators can also set up a hub such as a data center or a head end router. Doing that involves creating a hub profile, entering configuration parameters for that hub, and applying it to the router (Figure 22).

In terms of configuration, hub profile templates are very similar to WAN Edge Templates. Many of the same parameters such as general information, NTP servers, DNS settings, WAN and LAN definitions, traffic steering policies and application policies, and routing, are present in these types of templates.

The hub profiles will define overlays, which are essentially the SD-WAN topology and which will be used to configure

branches. Between the nodes in these overlays, traffic can be forwarded using **secure vector routing (SVR)**, which provides the deny-by-default access, tunnel-free adaptive encryption, and 30-50% bandwidth savings compared to other SD-WAN implementations.

Day 1: Zero Touch Deployment

Having completed Day 0 operations, you can begin the process of deploying your devices using Zero Touch Provisioning (ZTP). These operations are referred to as Day 1 tasks.

Under Inventory (Figure 23), you can claim your devices, which (assuming they are assigned to a site) will be automatically configured via their respective templates.

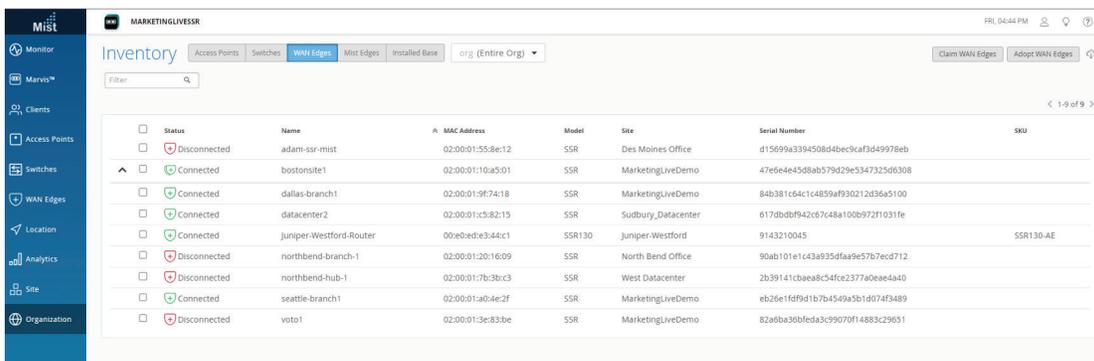


Figure 23: Claiming a WAN Edge Device (the SSR)

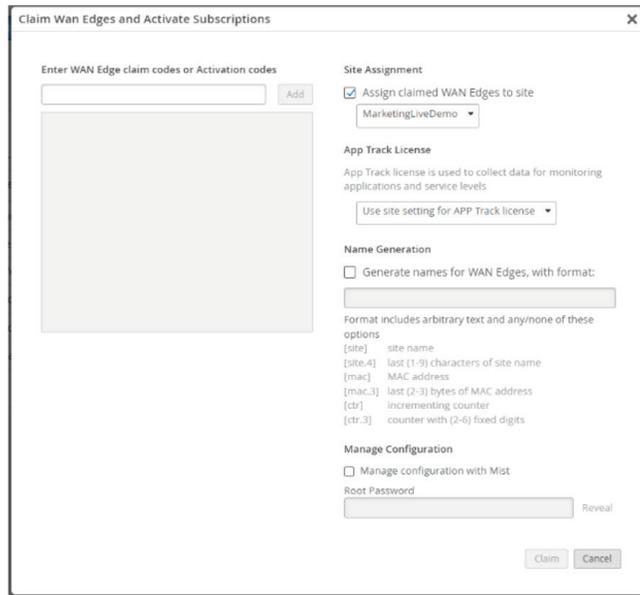


Figure 24: Entering the Claim Code(s) for SSRs

After selecting Claim WAN Edges, enter the claim code for the SSR, which can be found on the back of your device (Figure 24).

If you are claiming multiple devices, you can do so here. You must also enter the site assignment. Site variables can be called upon here.

You can enter license tracking information, and generate a name for the WAN edge devices; you can also enter the password to manage the devices using Mist.

You can claim the switches and access points in the same way (Figure 25).

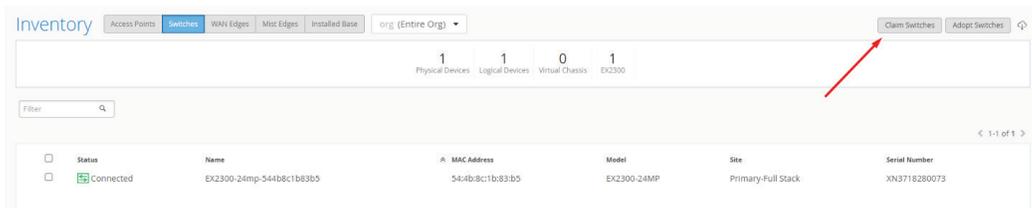


Figure 25: Claiming Switches

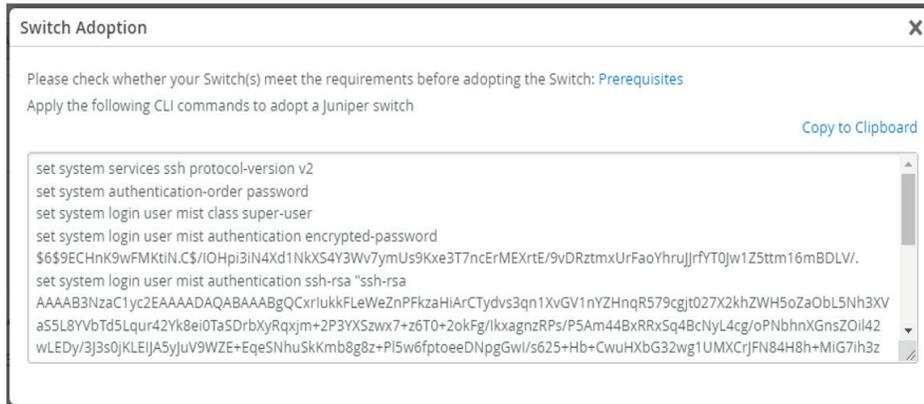


Figure 26: Adopting a Switch

For any device that doesn't have a claim code, you can "adopt" the device using a provided set of CLI commands (Figure 26).

As with all devices (WAN, Switch, or AP), you can use the [Mist AI Mobile App](#) from an application store (e.g., from [Apple](#) or [Google](#)) to get your device into inventory. You start this process by scanning the device's QR code with the mobile application (Figure 27).

After scanning the code, you then simply refresh the Inventory page for the related domain to see it appear.

Following that, the devices—as they connect to the Mist Cloud—will get their configurations assigned through the templates.

Once all of the device interfaces and their endpoints are defined (and templates assigned to sites and equipment deployed), the network will start to ensure connectivity by sending bidirectional forwarding detection (BFD) traffic between the endpoints. BFD will help measure jitter, loss, and latency along the path.



Figure 27: QR Code on an Access Point

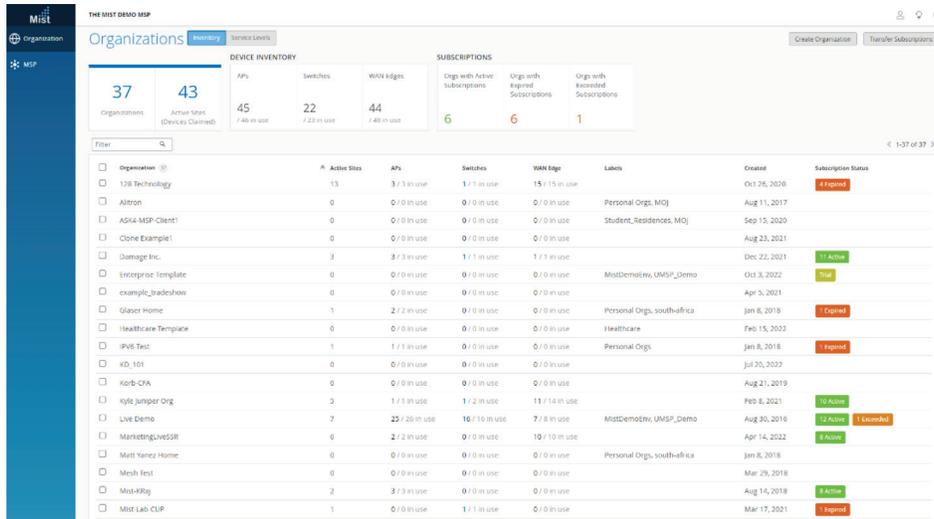


Figure 28: MSP Dashboard

Days 2 and Beyond: AI-driven Operational Benefits

Day 2 operations refer to all administrative and operational tasks after deployment. This includes troubleshooting, maintenance, and day-to-day operations.

The MSP Dashboard

As a Managed Service Provider (MSP) serving many enterprises, your task is simplified with the **Juniper MSP Dashboard**. You can manage your entire customer estate and quickly onboard new customers with correct attributes and configurations.

The dashboard (Figure 28) provides a view of all the organizations being managed.

For each organization, there is a list of all access points, switches, and WAN edges, along with the status of any subscriptions.

In a similar way that the MSP dashboard shows you all the organizations you manage, you can (by selecting an individual organization and using a drop-down menu) view all the sites within an organization (Figure 29).

Service Level Experiences (SLE)

For all of these organizations and sites (and in every network domain), you can manage **Service Level Experiences (SLE)** for users and operators. SLEs are maintained using applied data science and machine learning to understand the actual end user experience on the network.

Going well beyond the concept of client health (i.e., whether a device is functioning or not), SLEs measure actual user experiences in real-time, based on continuously delivered (and actionable) telemetry.

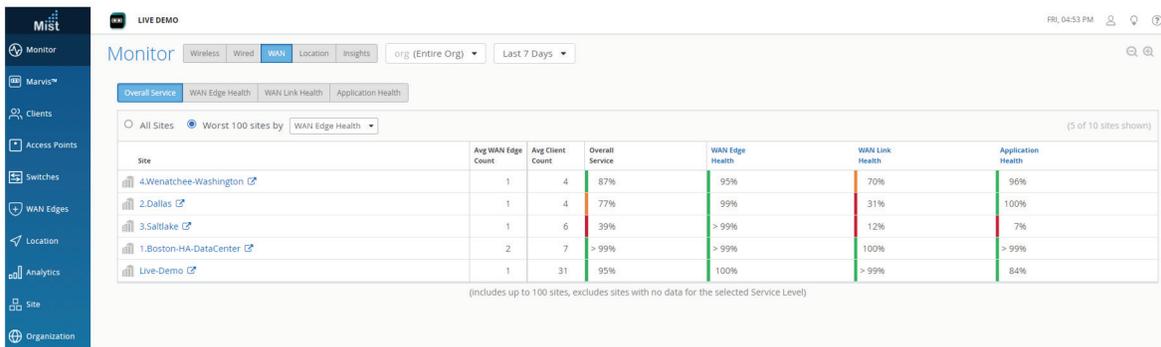


Figure 29: Status of Sites Within an Organization

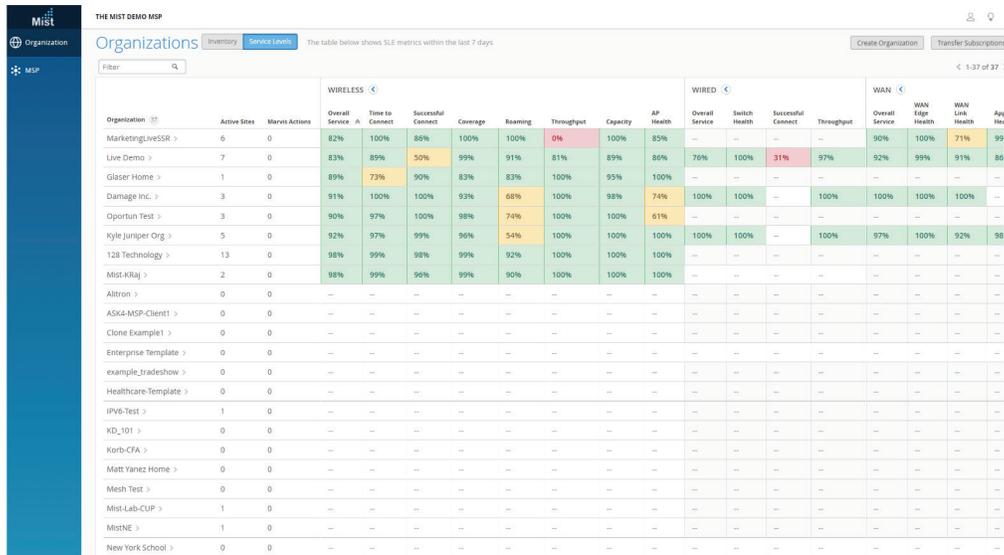


Figure 30: SLEs by Organization

As an MSP, you can check SLEs for all the organizations being managed (Figure 30).

The SLEs are color-coded so that serious problems appear in yellow or red. This lets you correlate problems across domains, which is discussed in the following sections.

SLEs by Domain

The ability to view SLEs by domain (wireless, wired and WAN), and to correlate them, is extremely powerful. The different domains (Table 1) are interrelated: problems in each domain often affect other domains.

Table 1. SLEs by Domain

Wireless	Wired	WAN
Overall Service	Overall Service	Overall Service
Time to Connect	Switch Health	WAN Edge Health
Successful Connections	Successful Connections	WAN Link Health
Coverage	Throughput	Application Health
Roaming		
Throughput		
Capacity		
Health		

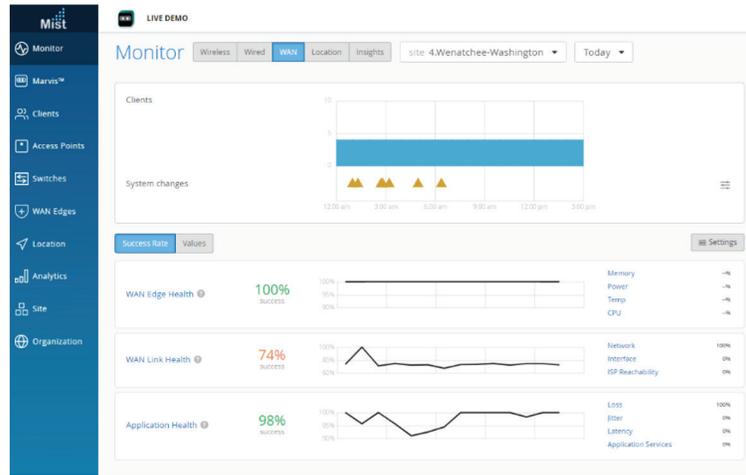


Figure 31: WAN Statistics

Other debugging sessions may begin with SLEs in any domain. Toggling to WAN statistics, for instance, you can view SLEs for WAN Edge health, WAN Link health, or application health (Figure 31).

For each device, you can see upstream paths or applications that are being accessed.

Similarly, clicking the Wired button, you can check the relevant characteristics such as throughput, successful connections, or switch health (Figure 32).

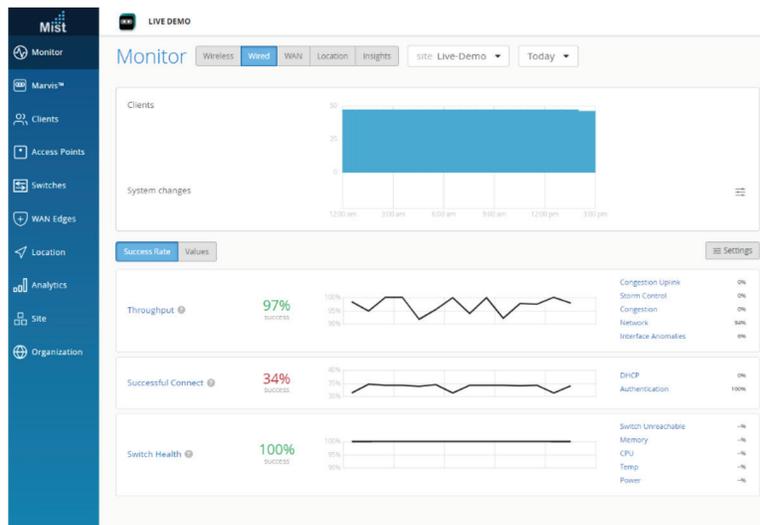


Figure 32: Wired Switching Statistics

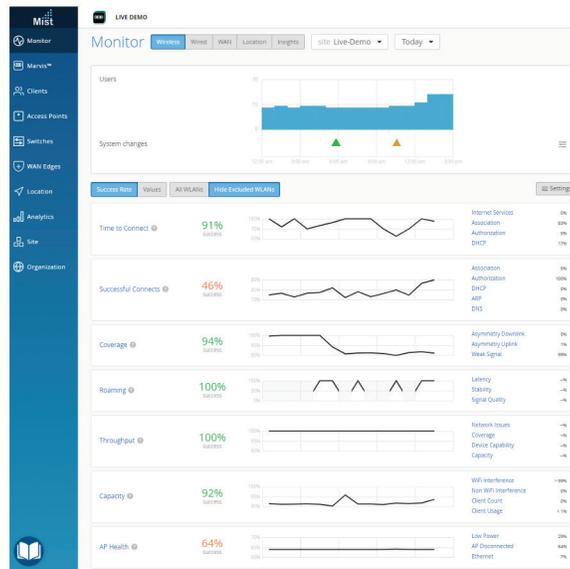


Figure 33: Wireless Statistics

For wireless devices, you can check, for example, coverage, roaming, capacity, and overall AP health (Figure 33).

Using SLE Classifiers

To help you find the root causes of problems, SLEs are subdivided into classifiers. SLE classifiers may include

authentication or authorization, latency, signal strength, and other characteristics that help you diagnose and correct issues.

For instance, if you see unsuccessful connections, you can determine if the problem is one of authorization (Figure 34).

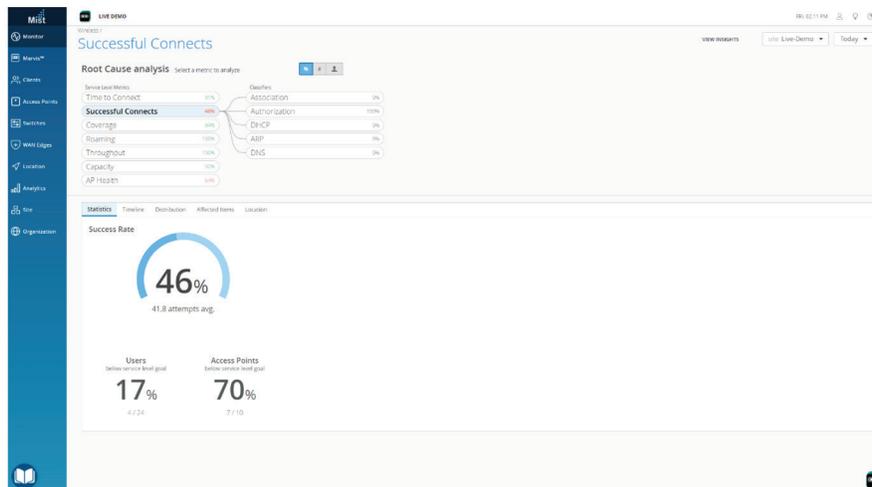


Figure 34: Use of Classifiers to Isolate a Problem

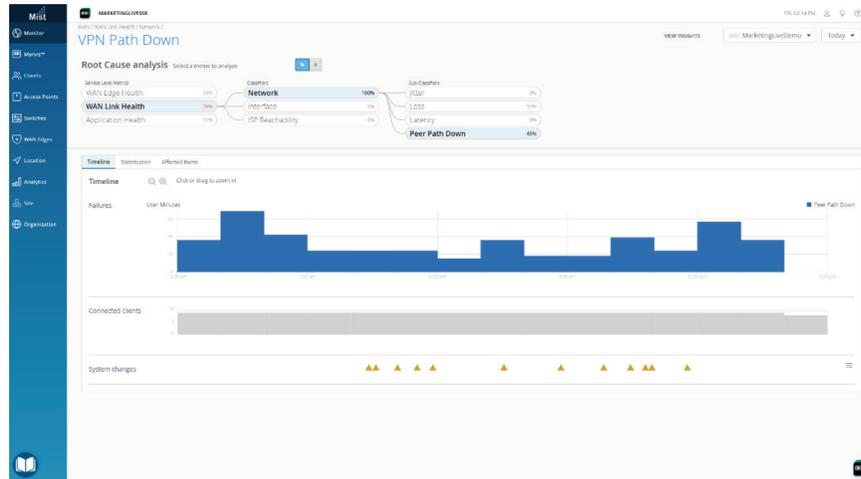


Figure 35: Diagnosis of a VPN Path Problem

You can see which SSRs, paths, and clients are associated with each site, and isolate the causes of problems such as an intersite VPN path being down (Figure 35).

Using the guidance provided by SLEs and their classifiers, you can gather further information on relevant events by viewing the Monitor page under an organization (Figure 36).

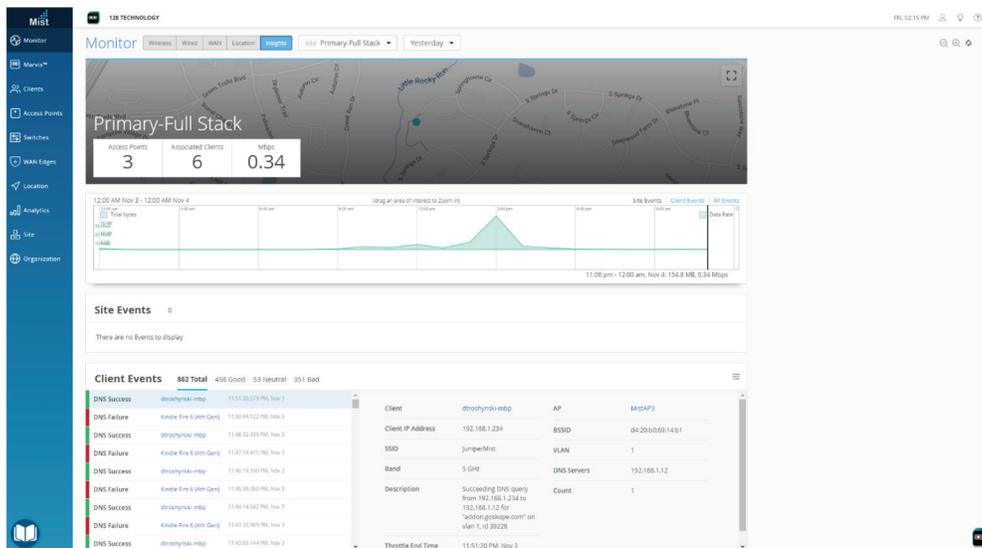


Figure 36: Monitor Page Under an Organization

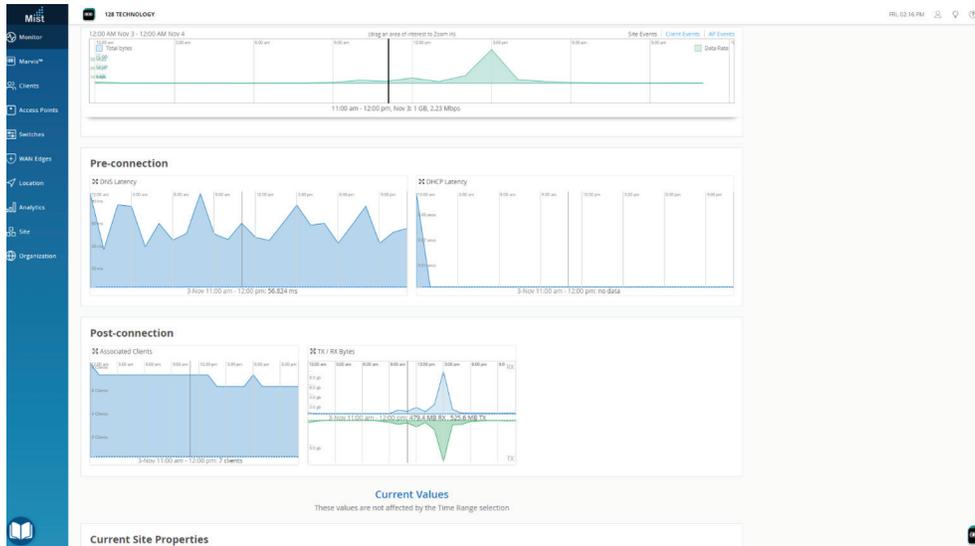


Figure 37: Throughput Shown from the Monitor Page

Under the Insights view, this page provides an overview of what each device is doing and any events that are occurring. You can see site events, client events, and access point events. Monitor also helps you see (at a glance) which applications people are accessing on the network. You can also see what type of traffic is passing, along with bandwidth usage, throughput, and connection efficiency (Figure 37).

In other per-domain diagnostics, you can see which ports are up or down, and look at throughput, connections, and switch health. For access points, you can see which ports are down at a particular site, along with which clients are attached to which switches. **Indoor Location Services** helps you determine whether some access points are overloaded with too many users, and can also help with site planning.

This level of visibility, event correlation and problem isolation is state of the art, and the AIDE goes further with AI/ML capabilities that will lead quickly to problem resolution, and can in many cases correct problems in the network for you.

The Marvis Virtual Network Assistant

The **Marvis Virtual Network Assistant (VNA)** recognizes behavior on the network and offers suggestions as to how problems can be fixed across the Wired, Wireless, and WAN domains.

Marvis Actions, for instance, divides into different possible problems for each device type (Figure 38).

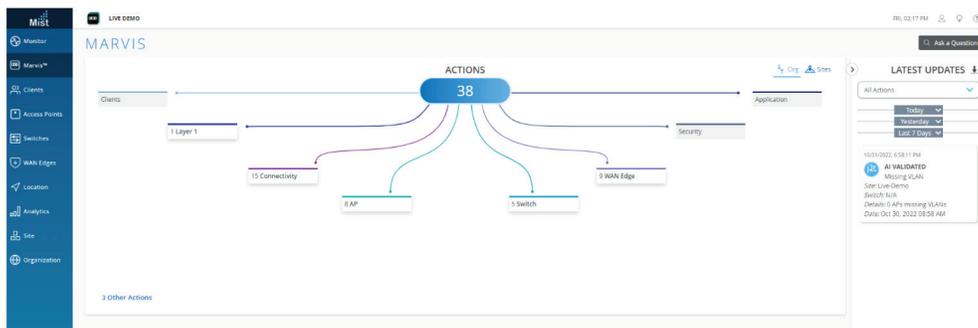


Figure 38: Marvis Actions

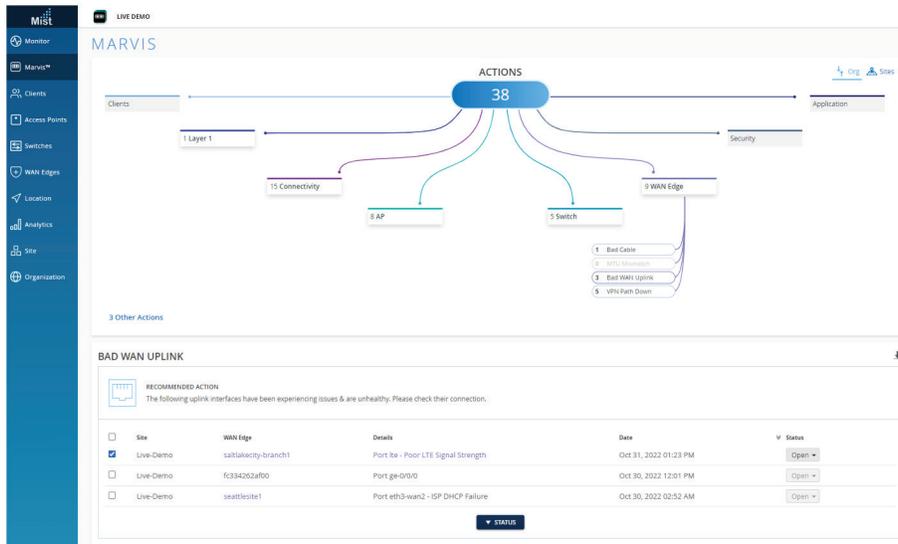


Figure 39: Marvis Diagnosis and Suggestion at the WAN Edge

With each of these options, Marvis offers suggestions as to how to fix them. For instance, for a weak LTE signal, Marvis suggests that it could be an ISP problem with DHCP or simply caused by a bad cable (Figure 39).

Similarly, for switches, Marvis may indicate whether the cause of a problem is a missing VLAN or a bad cable. For access points, you can do a health check for connectivity, authorization failures, DHCP, ARP, or (again) a Layer 1 problem such as a bad cable.

With Marvis, you click on an Action and see where the problem lies. This represents a major breakthrough for troubleshooting capabilities, a greatly simplified and proactive process from traditional techniques.

Additionally, Marvis offers very specific help through a natural language processing (NLP) chat interface. You can type in specified search terms, user names, a trouble ticket number, or questions (Figure 40).

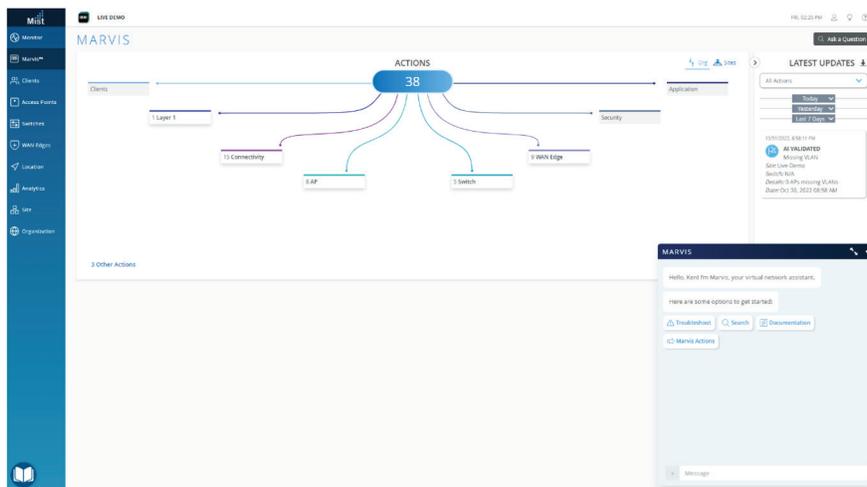


Figure 40: Marvis Natural Language Processing (Chat)

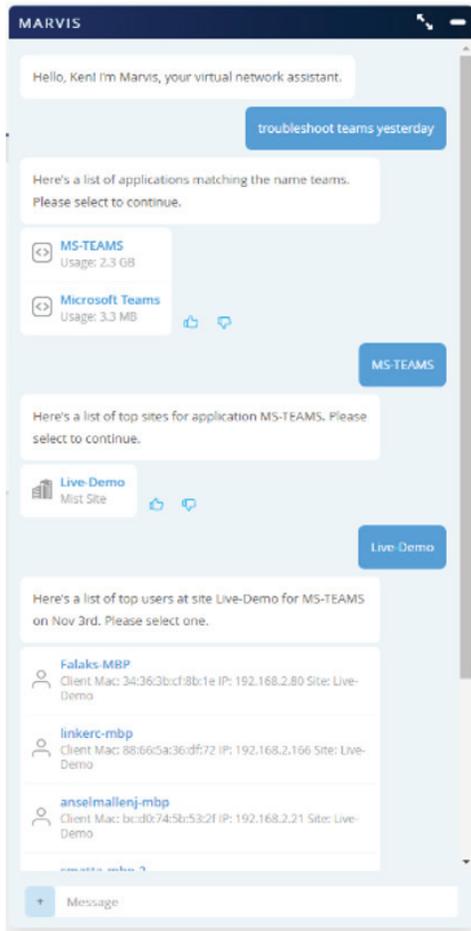


Figure 41: Asking Marvis to Troubleshoot Teams

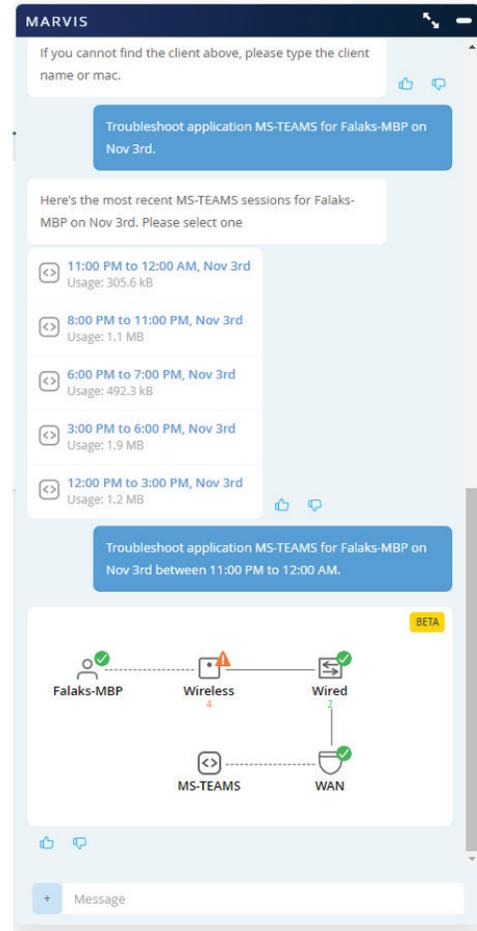


Figure 42: Graphical Diagnostics from Marvis

For example, let's say users are reporting trouble with Microsoft Teams. You can suggest different sites to look at, and can find users from Teams. Depending on how clients are set up, you may get an answer just by typing "troubleshoot teams" in the chat box (Figure 41).

As you scroll through suggestions, Marvis will show a simple graphic in the chat box to help you find where the problem is. From the user to an application—including all intermediate hops—Marvis helps you locate the source of the problem (Figure 42).

From the Marvis chat, you can quickly determine what is working and what is not, indicated by either a caution icon or green check mark. You may see that the switch, the router, and Teams is all working correctly, but the wireless is not. In a case like that, the problem may be in the access point itself.

This level of insight greatly reduces troubleshooting time. Furthermore, isolating by time period is very valuable; you may notice, for example, that in a different time slice it may be that the WAN router was the source of the issue.

Note: For additional ease of use, there are mobile applications for Mist Experience and the Marvis Client available for Android or Apple devices.

Conclusion

This solution guide has provided a summary of what can be accomplished with Juniper's AI-driven Enterprise in a full stack deployment. Simplicity is the key—the Mist AI Cloud is designed to be easy to use and to make the network easy to troubleshoot across all domains.

This is how you deploy, maintain and operate networks with an AI-driven Enterprise—it is the essence of experience-first networking for both users and administrators.

Resources

Mist Management and Live Demo

- [Juniper Mist Live Demo](#) at [Manage.mist.com](#)

Documentation

- [Guided Setup for Mist WAN Assurance](#)
- [Configure and Manage the EX Switch and Mist AP in the Juniper Mist Cloud](#)
- [Overview of the Cloud-Managed Midsize Branch Office](#)

Solution Briefs and White Papers

- [Client-to-Cloud Assurance with an AI-driven Enterprise](#)
- [Session Smart Routing: How it Works](#)

Data Sheets

- [SSR 100 Line of Routers](#)
- [WAN Assurance](#)
- [Wireless Assurance](#)
- [Wired Assurance](#)

Web Pages

- [Juniper Networks Presents at Network Field Day 23 \(July 2022\)](#)
- [Managed Service Provider \(MSP\) Dashboard](#)
- [Mist AI and Cloud](#)
- [Marvis Virtual Network Assistant](#)
- [Session Smart Routers \(SSR\)](#)
- [EX switches](#)

- [Indoor Location Services](#)
- [Mist access points](#)
- [Mist AI Mobile App](#)
- [Mist Experience Mobile App](#)
- [Marvis Client](#)
- [Service Level Experiences](#)
- [Juniper Unified Managed Services Program](#)
- [WAN Assurance](#)
- [Wired Assurance](#)
- [Wireless Assurance](#)

Videos

- [WAN Assurance](#)
- [Juniper Networks Wired and Wireless Integration](#)
- [Juniper Networks AI-driven SD-WAN](#)

About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily.

At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business.



Driven by
Experience™

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.207.125.700
Fax: +31.207.125.701

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000 | Fax: +1.408.745.2100
[www.juniper.net](#)

